GENERAL

Special Issue: RV 2021

Monitoring with verified guarantees

Jan Baumeister¹ · Johann C. Dauer² · Bernd Finkbeiner¹ · Sebastian Schirmer²

Accepted: 30 May 2023 / Published online: 11 July 2023 © The Author(s) 2023

Abstract



Runtime monitoring is generally considered a light-weight alternative to formal verification. In safety-critical systems, however, the monitor itself is a critical component. For example, if the monitor is responsible for initiating emergency protocols, as proposed in a recent aviation standard, then the safety of the entire system critically depends on the correctness of the monitor. In this paper, we present a verification extension to the LOLA monitoring language that extends the efficient specification of the monitor with Hoare-style annotations that guarantee the correctness of the monitor specification. We add two new operators, assume and assert, which specify assumptions of the monitor and expectations on its output, respectively. The validity of the annotations is established by an integrated SMT solver. We report on experience in applying the approach to specifications from the avionics domain, where the annotation with assumptions and assertions has lead to the discovery of safety-critical errors in specifications. The errors range from incorrect default values in offset computations to complex algorithmic errors that result in unexpected temporal patterns. We also report how verified specifications can be monitored efficiently at runtime.

Keywords Formal methods · Cyber-physical systems · Runtime verification · Hoare logic

1 Introduction

Cyber-physical systems are inherently safety-critical due to their direct interaction with the physical environment – failures are unacceptable. A means of protection against failures is the integration of reliable monitoring capabilities. A *monitor* is a system component that has access to a wide range of system information, e.g., sensor readings and control decisions. When the monitor detects a failure, i.e., a violation of the behavior stated in its *specification*, it notifies the system or activates recoveries to prevent failure propagation.

The task of the monitor is critical to the safety of the system, and its correctness is therefore of utmost importance.

 S. Schirmer sebastian.schirmer@dlr.de
 J. Baumeister jan.baumeister@cispa.de
 J.C. Dauer johann.dauer@dlr.de
 B. Finkbeiner finkbeiner@cispa.de

¹ Helmholtz Center for Information Security (CISPA), Saarbrücken, Germany

² German Aerospace Center (DLR), Braunschweig, Germany

Runtime monitoring approaches like LOLA [8, 11] address this by describing the monitor in a formal specification language, and then generating a monitor implementation that is provably correct and has strong runtime guarantees, for example, on memory consumption. Formal monitoring languages typically feature temporal [24] and sometimes spatial [21] operators that simplify the specification of complex monitoring behaviors. However, the specification itself, the central part of runtime monitoring, is still prone to human errors during specification development. Hence, how can we check that the monitor specification itself is correct?

In this paper, we introduce a verification feature to the LOLA framework. Specifically, we extend the specification language with *assumptions* and *assertions*. The framework statically verifies that the assertions are guaranteed to hold if the input to the monitor satisfies the assumptions. This verification feature was previously introduced in [9]. Here, we extend this work by providing a proof of soundness and presenting an online monitoring approach with experimental results that checks the satisfaction of assumptions during runtime to activate assertion checks.

The prime application area of LOLA is unmanned aviation. LOLA is increasingly used for the development and operation monitoring of unmanned aircraft; for example, the LOLA monitoring framework has been integrated into the



Fig. 1 Run-Time Assurance architecture proposed by ASTM F3269-21 to safely bound a complex function using a safety monitor

DLR unmanned aircraft superARTIS¹ [3] using an FPGA realization [1]. The verification extension presented in this paper is motivated by this work. In practice, system engineers report that support for specification development is necessary, e.g., sanity checks and proofs of correctness. Additionally, recent developments in unmanned aviation regulations and standards indicate a similar necessity. One such development is the industry standard F3269-21 (Standard Practice for Methods to Safely Bound Flight Behavior of Unmanned Aircraft Systems Containing Complex Functions) by ASTM International.² ASTM F3269-21 introduces a certification strategy based on a Run-Time Assurance (RTA) architecture that bounds the behavior of a complex function by a safety monitor [20], similar to the well-known Simplex architecture [27]. This complex function could be a Deep Neural Network as proposed in [7]. A simplified version of the architecture³ of ASTM F3269-21 is shown in Fig. 1.

At the core of the architecture is a safety monitor that takes the inputs and outputs of the complex function, and decides whether the complex function behaves as expected. If not, the monitor switches the control from the complex function to a matching recovery function. For instance, the flight of an unmanned aircraft could be separated into different phases, e.g., take-off, cruise flight, and landing. For each of these phases, a dedicated recovery could be defined, e.g., braking during take-off, the activation of a parachute during cruise flight, or a go-around maneuver during landing. Further, it is crucial that recoveries are only activated under certain conditions and that only one recovery is activated at a time. For instance, a parachute activation during a landing approach is considered safety-critical. The verification extension of LOLA introduced in this paper can be used to guarantee statically that such decisions are avoided within the monitor specification. Consider the simplified LOLA specification

² https://www.astm.org/

```
^{3}\, In its original version, data is separated into assured and unassured data and data preparation components are added.
```

```
input event_a, event_b, value: Bool, Bool, Float32
assume <a1> !(event_a and event_b)
output braking : Bool := ...computation...
output parachute : Bool := ...computation...
output go_around : Bool := ...computation...
assert <a1> !(braking and parachute)
```

that declares an assumption on the system input events and asserts that braking and parachute never evaluate to *true* simultaneously.

In the following, we first give a brief introduction to the stream-based specification language LOLA, then present the verification approach, and give details on the tool implementation and our tool experience with specifications that were written based on interviews with aviation experts. Last, we consider the case where assumptions might not be satisfied during runtime. Our results show that standard LOLA specifications are indeed prone to error, and that these errors can be caught with the formal verification introduced by our extension.

Related work Most work on the verification of monitors focuses on the correct transformation into a general programming language. For example, Copilot [22] specifications can be compiled into C code with constant time and memory requirements. Similarly, there is a translation validation toolkit for LOLA monitors implemented in Rust [11], which is based on the Viper verification tool [19]. Translation validation of this type is orthogonal to the verification approach of this paper. Instead of verifying the correctness of a transformation, our focus is to verify the specification itself. Both activities complement each other and facilitate safer future cyber-physical systems.

Our verification approach is based on classic ideas of inductive program verification [12, 16], and is closely related to the techniques used in static program verifiers like KEY [4], Why3 [6], and Dafny [18]. In a verification approach like Dafny, we are interested in functional properties of procedures, specified as post-conditions that relate the values upon the termination of the procedure with those at the time of entry to the procedure, e.g., *ensure* y = old(y). By contrast, a stream-based language like LOLA allows arbitrary access to past and future stream values. This makes it necessary to *unfold* the LOLA specification in order to properly relate the assumptions and assertions in time.

Most closely related to stream-based monitoring languages are synchronous programming languages like LUS-TRE [15], ESTEREL [5], and SIGNAL [13]. For these languages, the compiler is typically used for verification – a program representing the negation of desired properties is compiled with the target program and a check for emptiness decides whether the properties are satisfied. Furthermore, a translation from past linear-time temporal logic to ES-TEREL was proposed to simplify the specification of more

¹ https://www.dlr.de/content/en/research-facilities/superartis-en. html

complex temporal properties [17]. Other verification techniques also exist like SMT-based *k*-Induction for LUS-TRE [14] or a term rewriting system on synced effects [28]. A key difference in our approach is that we do not rely on compilation. Our verification works at the level of an intermediate representation. Furthermore, synchronous programming languages are limited to past references, while the stream unfolding for the inductive correctness proof of the LOLA specification includes both past and future temporal operators. Similar to *k*-Induction, our approach is sound but not complete.

2 Runtime monitoring with Lola

LOLA is a stream-based language that describes the translation from input to output streams:

```
input t_1 : T_1

:

input t_m : T_m

output s_1 : T_{m+1} := e_1(t_1, \dots, t_m, s_1, \dots, s_n)

:

output s_n : T_{m+n} := e_n(t_1, \dots, t_m, s_1, \dots, s_n)

trigger \varphi message
```

where input streams carry synchronous arriving data from the system under scrutiny, output streams represent calculations, and triggers generate notification *messages* at instants where their condition φ becomes *true*. Input streams t_1, \ldots, t_m and output streams s_1, \ldots, s_n are called *independent* and *dependent variables*, respectively. Each variable is typed: independent variables t_i are typed T_i and dependent variables s_i are typed T_{m+i} . Dependent variables are computed based on *stream expressions* e_1, \ldots, e_n over dependent and independent stream variables. A (stream) expression is one of the following:

- an atomic expression *c* of type *T* if *c* is a constant of type *T*;
- an atomic expression *s* of type *T* if *s* is a stream variable of type *T*;
- an expression *ite*(*b*, *e*₁, *e*₂) of type *T* if *b* is a Boolean expression and *e*₁, *e*₂ are expressions of type *T*. Note that *ite* abbreviates the control construct *if-then-else*;
- an expression $f(e_1, \ldots, e_k)$ of type T if $f: T_1 \times \cdots \times T_k \mapsto T$ is a k-ary operator and e_1, \ldots, e_k are expressions of type T_1, \ldots, T_k ;

• an expression *o.offset*(*by* : *i*).*defaults*(*to* : *d*) of type *T* if *o* is a stream variable of type *T*, *i* is an Integer, and *d* is of type *T*.

For example, consider the LOLA specification

```
input altitude: Float32 // in m
output altitude_bound := altitude > 200.0
trigger altitude_bound "Warning: Decrease altitude!"
```

that notifies the system if the current altitude is above its operating limits, i.e., 200.0 meters. Note that stream types are inferred, i.e., altitude_bound is of type Bool.

LOLA uses temporal operators that allow output streams to access its and others previous and future stream values. The stream

```
output alt_count := if altitude ≤ 200.0 then 0
    else alt_count.offset(by: -1).defaults(to: 0) + 1
```

represents a count of consecutive altitude violations by accessing its own previous value, i.e., offset(by: x) where a negative and positive integer x represents past and future stream accesses, respectively. Since temporal accesses are not always guaranteed to exist, the default operator defines values which are used instead, i.e., defaults(to: d) where d has to be of the same type as the used stream. Here, at the first position of alt_count the default value zero is taken. As abbreviations for the temporal operators, alt_count[x, d] is used. Further, $s[x..y, d, \circ]$ for x < y abbreviates $s[x,d] \circ s[x+1,d] \circ \cdots \circ s[y,d]$ where \circ is a binary operator. Using alt_count > 10 as a trigger condition is preferable if only persistent violations should be reported.

In general, LOLA is a specification language that allows to specify complex temporal properties in a precise, concise, and less error-prone way. The focus is on *what* properties should be monitored instead of *how* a monitor should be executed. Therefore, the LOLA monitor synthesis automatically infers and optimizes implementation details like evaluation order and memory management. The evaluation order [11] of LOLA streams is automatically derived by analysis of the *dependency graph* [8] of the specification. This allows ignoring the order when taking advantage of the modular structure of LOLA output streams, e.g.,

```
output alt_avg := alt_count / (position+1)
output alt_count := if altitude ≤ 200.0 then 0
    else alt_count.offset(by: -1).defaults(to: 0) + 1
output pos := pos.offset(by: -1).defaults(to: 0)
```

where pos and alt_count are used before their definition. Further, the graph allows to detect all invalid cyclic stream dependencies, e.g.,

output a := a.offset(by:0).defaults(to:0) .

3 Assumptions and assertions

In this section, we present the verification extension for the LOLA specification language. Note that a color version of this paper is given online. The extension allows the developer to annotate the LOLA specification with *assumptions* and *assertions* in order to verify the desired guarantees on the computed streams. As an example, consider the simplified specification in Listing 1, which is structured into stream computations in Lines 1 to 28, and assumptions and assertions from Line 30 onwards.

The computation part specifies a safety monitor within a RTA architecture that triggers recovery functions for three different flight phases. First, the take-off recovery function is triggered (Line 24) when the targeted take-off speed was not achieved on a runway up to a predefined point (Lines 14–15). The distance between the current position and the end of the runway with local coordinates (0,0) is computed in Line 9. Second, in-flight a parachute is activated (Line 26) when virtual barriers for the aircraft, i.e., a geofence, are exceeded (Line 17) [26]. Last, during landing, up to a point of no return (alt < 10.0), a new landing attempt is initiated (Line 27)

```
input alt : Float32 // Height above ground
                                                               1
                                                               2
input x, y : Float32, Float32 // Position
input speed : Float32 // Velocity of aircraft
                                                               3
input landing : Bool // Indicates landing mode
                                                               4
                                                               5
// Status of landing gear
input lg_status : (Float32,Float32,FLoat32)
                                                               6
                                                               7
// Complex computations
                                                               8
                          : Float32 := \sqrt{x^2+y^2}
                                                               9
output dst_on_runway
                                      := ...
                                                               10
output geofence_violation : Bool
output landing_gear_ready : Bool
                                       := ...
// Take-off contingency
                                                               13
output decelerate := alt < 1.0 \land speed < 10.0
                                                               14
                    \land dst_on_runway > 20.0
// In-flight contingency
<code>output</code> parachute := geofence_violation \land alt > 100.0
                                                               17
// Landing contingency
                                                               18
output gain_alt := landing \land alt \ge 10.0
                                                               19
                \land (landing_gear_ready[-4..0, true, \land]
                                                               20
                 \rightarrow speed > 10.0)
                                                               21
                                                               22
                                                               23
// Notifications to the system
trigger decelerate "RECOVERY: Stop take-off by
                                                               24
                                                               25
                     decelerating aircraft."
                                                               26
trigger parachute
                    "RECOVERY: Activate parachute."
                    "RECOVERY: Gain altitude for next
                                                               27
trigger gain_alt
                     landing attempt."
                                                               28
                                                               29
                                                               30
assume <a1> \neg(landing \land geofence_violation)
assume <a1> abs(speed) \leq 80.0
                                                               31
                                                               32
// Only one contingency is activated at once.
                                                               33
assert <a1> \neg( (decelerate \land parachute)
                                                               34
              \lor (decelerate \land gain_alt)
                                                               35
              \vee (parachute \wedge gain_alt) )
                                                               36
// Parachute only activated 100 m above ground.
                                                               37
                                                               38
<code>assert <a2> parachute \rightarrow alt > 100.0</code>
```

Listing 1 A simplified Run-Time Assurance LOLA specification with three recovery functions for three different flight phases. Assumptions and assertions are used to show that only one recovery function is activated at once

if the aircraft's speed is too large or its landing gear is not yet ready. To be more robust, the current and the previous value of the landing_gear_ready is taken into account (Lines 19–21).

With the verification extension, the specification assures that recoveries are not activated simultaneously (Lines 34–36), for instance, there is no possibility that a parachute is activated during a landing approach. The first two conjunctions in Lines 34 and 35 evaluate to false because relevant outputs use a disjoint altitude condition. The last conjunction requires an assumption. Here, two assumptions are linked by the identifier a1 to the assertion. The assumptions specify the known bound of received speed data (Line 31) as well as operational information (Line 30), e.g., given by the concept of operation a nominal landing is only foreseen within the predefined operational airspace. Note that assumptions are provided by the user and are assumed to be valid. Further, a second assertion is stated in Line 38 that guarantees that the parachute should only be activated when the aircraft is 100 meters above ground. In this case, the property can be shown assumption-free. Assertions help engineers show that certain properties are true. The given assertions indicate how specification debugging and management can benefit from the extension - it avoids digging into complex computations.

The extension and its verification approach are presented in the following. In general, the verification extension is used if a LOLA specification is annotated in the following way:

```
assume \langle \alpha_1 \rangle \quad \theta_1

\vdots

assume \langle \alpha_m \rangle \quad \theta_m

assert \langle \alpha_{m+1} \rangle \quad \psi_1

\vdots

assert \langle \alpha_{m+n} \rangle \quad \psi_n
```

where $\alpha_1, \ldots, \alpha_{m+n} \in \Gamma$ are identifiers for $\theta_1, \ldots, \theta_m, \psi_1, \ldots, \psi_n$, which are Boolean stream expressions with possibly temporal operators. For convenience, we define functions which return all θ and ψ that are linked to a given α identifier: $assume(\alpha) = \{\theta_j \mid \forall \alpha_j \in \Gamma, \alpha = \alpha_j\}$ and $assert(\alpha) = \{\psi_j \mid \forall \alpha_j \in \Gamma, \alpha = \alpha_j\}$. The set of assertions ψ_1, \ldots, ψ_n is *correct* for all input streams if and only if whenever an assumption is satisfied, its corresponding assertion is satisfied as well.

The verification of assertions relies on the encoding of the LOLA execution in Satisfiability Modulo Theory (SMT). We define the *smt* function that encodes a stream expression

in Definition 1. It can be used to encode independent and dependent variables as well as expressions of assumptions and assertions.

Definition 1 (SMT-Encoding of Stream Expressions)

Let Φ be a LOLA specification over independent stream variables t_1, \ldots, t_m and dependent stream variables s_1, \ldots, s_n . Further, let the natural number N + 1 be the length of the input streams, c be an SMT constant symbol, and $\tau_1^0, \ldots, \tau_1^N, \ldots,$ $\tau_m^0, \ldots, \tau_m^N, \sigma_1^0, \ldots, \sigma_1^N, \ldots, \sigma_n^0, \ldots, \sigma_n^N$ be SMT variables. Then, the function *smt* recursively encodes a stream expression e at position j with $0 \le j \le N$ in the following way:

- Base cases:
 - smt(c)(j) = c,
 - $smt(t_i)(j) = \tau_i^j,$
 - $smt(s_i)(j) = \sigma_i^j;$
- Recursive cases:

$$- smt(f(e_1,...,e_n))(j) = f(smt(e_1)(j),..., smt(e_n)(j)), - smt(ite(e_b,e_1,e_2))(j) = ite(smt(e_b)(j), smt(e_1)(j), smt(e_2)(j)), - smt(e[k,c])(j) = {smt(e)(j+k) if 0 \le j+k \le N, c otherwise, }$$

where ite is an SMT encoding of *if-then-else*; f is an interpreted function if f is from a theory supported by the SMT solver and an uninterpreted function otherwise.

Next, Proposition 1 establishes the correctness of asserted stream properties for finite input streams. If the set of assertions is correct, asserted stream properties are guaranteed to be valid in each step of the monitor execution. In practice, such specifications are preferable. In the following, let Φ be a LOLA specification with verification annotations. Further, we refer to the set of input streams and computed output streams as stream execution.

Proposition 1 (Assertion Verification of a Finite Stream Execution)

Let Φ be a LOLA specification and let s_1, \ldots, s_n be dependent stream variables used in Φ . The set of assertions is correct for a finite stream execution with length N + 1 under given assumptions, if the following formula is valid:

$$\bigwedge_{0 \le i \le N} \Big(\bigwedge_{\alpha \in \Gamma} \Big(\bigwedge_{\theta \in assume(\alpha)} smt(\theta)(i) \land$$
$$\bigwedge_{s_k \in \Phi} \sigma_k^i = smt(e_k)(i) \to \bigwedge_{\psi \in assert(\alpha)} smt(\psi)(i) \Big) \Big).$$

The formula in Proposition 1 unfolds the complete stream execution and informally expresses that an assertion must

hold in each stream position whenever its corresponding assumption and implementation are satisfied.

To avoid the complete unfolding and allow arbitrary stream lengths, an inductive argument is given in Proposition 2 that defines proof obligations for an annotated LOLA specification. Next, we present a template for the stream unfolding that helps to define the proof obligation at the *Begin*ning (Definition 3), during *Run* (Definition 4), and at the *End* (Definition 5) of a stream execution.

Definition 2 (Template Stream Unfolding)

We define the template formula ϕ_t that states proof obligations as:

$$\bigwedge_{\alpha \in \Gamma} \left(\bigwedge_{i \in p_asm} \left(\bigwedge_{\substack{\theta \in assume(\alpha)}} smt(\theta)(i) \right) \right)$$
$$\land \bigwedge_{i \in p_asserted} \left(\bigwedge_{\substack{\psi \in assert(\alpha)}} smt(\psi)(i) \right)$$
$$\land \bigwedge_{i \in p_streams} \left(\bigwedge_{\substack{0 < k \le n}} \sigma_k = smt(e_k)(i) \right)$$
$$\rightarrow \bigwedge_{i \in p_assert} \left(\bigwedge_{\substack{\psi \in assert(\alpha)}} smt(\psi)(i) \right) \right),$$

where *p_asm*, *p_asserted*, *p_streams*, and *p_assert* are template parameters. They are sets of positions for the unfolding of assumptions, previously proven assertions, output streams, and assertions, respectively.

The template formula in Definition 2 uses template parameters for the stream unfolding. For instance, the parameter assignment $p_asm := \{i \mid 0 \le i < 10\}$ adds assumptions at the first ten positions of the stream execution. Further, the parameter $p_asserted$ allows incorporating the induction hypothesis.

In the following, we will use the LOLA specification in Listing 2 as a running example for the template stream unfolding.

Here, the input *reset* represents a reset command for the output stream ol that counts how long no *reset* occurred. Output ol is used by output o2 which aggregates over the previous, the current, and the next outcome of ol. This is achieved by using the offset operator, e.g., o1[-1, 0] accesses the previous value of o1 if it exists, otherwise it takes the default value 0. As assertion, we show that o2 is always positive and never larger than three given the assumption that in each execution step either the previous or the next *reset* is *true*. The assumption ensures that at most two consecutive *resets* are *false*. Given the *reset* sequence of input values $\langle true; false; false \rangle$ that satisfies the assumption, the resulting ol stream evaluates to $\langle 0; 1; 2 \rangle$. Here, at the second position of the sequence, o2 evaluates to three. To show that the

```
assume<al> reset[-1, false] \v reset[1, false]
input reset : Bool
output o1 := if reset then 0 else o1[-1, 0] + 1
output o2 := o1[-1, 0] + o1 + o1[1, 0]
assert<al> 0 ≤ o2 and o2 ≤ 3
```

Listing 2 LOLA specification with assumptions on a reset that guarantees that an output remains within bounds

assertion also holds at the first and the last position of the sequence, out-of-bounds values must be considered.

We show how the template ϕ_t can be used at the beginning of a stream execution. Here, default values due to past stream accesses beyond the beginning of a stream need to be captured by the obligation to guarantee that the assertions hold in these cases. The combination of past out-of-bounds and future out-of-bounds default values must also be covered by the obligations in case the stream is stopped early. These scenarios are depicted for the running example in Fig. 2.

The figure shows four finite stream executions with different lengths. All stream positions are colored gray, while only some positions contain a single red dot. These features indicate the unfolding of stream variables and annotations using the template ϕ_t . A gray-colored position means that the assumptions have been unfolded and a dotted position means the assertion has been unfolded. Further, arrows indicate temporal stream accesses where solid lines correspond to accesses by outputs and dashed lines correspond to accesses by annotations, i.e., assumptions and assertions. For each stream execution, only the arrows for a single position are depicted - the arrows for other positions have been omitted for the sake of clarity. For example, for N = 0, the accesses of output o2 are both out-of-bounds, i.e., the default value zero is used. While for N = 3, the accesses at the second position are shown where only the past access of the assumption leads to an out-of-bounds access, i.e., only the dotted line towards the beginning of the stream execution. The figure depicts all necessary stream executions that cover all combinations of past out-of-bounds accesses, i.e., with and without future bound violations. The described unfoldings of Fig. 2 are formalized as proof obligations in Definition 3.

Definition 3 (Proof Obligations ϕ_{Begin} for Past Out-ofbounds Accesses)

Let $w_p = \sup(\{0\} \cup \{|k| \mid e[k,c] \in \Phi \text{ where } k < 0\})$ be the most negative offset and $w_f = \sup(\{0\} \cup \{k \mid e[k,c] \in \Phi \text{ where } k > 0\})$ be the greatest positive offset. The proof obligations ϕ_{Begin} for past out-of-bounds accesses are defined as the conjunction of template formulas:





Fig. 2 Four stream executions of different length N + 1 with the respective template unfolding are depicted. The stream executions consider all cases with past out-of-bound accesses. A gray-colored box indicates that an assumption has been unfolded at this position, while a red dotted box indicates that an assertion has been unfolded at this position. Solid and dashed arrows indicate accesses by streams and annotations, respectively

with template parameters:

2

3

4

5

 $\begin{array}{ll} \bullet \ p_asm & := \{i \mid 0 \le i \le N\}, \\ \bullet \ p_asserted & := \emptyset, \\ \bullet \ p_streams & := \{i \mid 0 \le i \le N\}, \\ \bullet \ p_assert & := \{i \mid 0 \le i < max(1, \min(N+1, 2 \cdot w_p))\}. \end{array}$

Next, the case where no out-of-bounds access occurs is considered. Hence, the obligations capture the nominal case where no default value is used. Since we have shown that past out-of-bounds accesses are valid we can use these proven assertions as assumptions. Figure 3 depicts a stream execution with a single dotted position, i.e., the position where the assertion must be proven. As can be seen, all accesses from this position are within bounds. Further, note that the accesses of the first and the last unfolded assumption, i.e., the first and the last gray-colored position, are also within bounds. The described unfolding is formalized as proof obligations in Definition 4.

Definition 4 (Proof Obligations ϕ_{Run} for No Out-ofbounds Accesses)

The proof obligations ϕ_{Run} without out-of-bounds accesses are defined as $\phi_t(p_asm, p_asserted, p_streams, p_assert)$ with template parameters:

• p_asm	$:= \{ i \mid w_p \le i \le N - w_f \},\$
• p_asserted	$:= \{i \mid 2 \cdot w_p \le i \le N - 2 \cdot w_f$
	$\wedge i \neq 3 \cdot w_p$ },
• p_streams	$:= \{i \mid 2 \cdot w_p \le i \le N - 2 \cdot w_f,\$
• p_assert	$:= \{i \mid i = 3 \cdot w_p\},\$

where $N = 3 \cdot (w_p + w_f)$.

Last, we consider the case where only future out-ofbounds accesses occur. Hence, the respective obligations need to incorporate default values of future out-of-bounds accesses. As before, we can use the previously proven assertions as assumptions. Figure 4 depicts a stream execution



Fig. 3 A stream execution of length N + 1 with the corresponding template unfolding is depicted. The stream execution considers the case where no out-of-bound access occurs. Gray-colored and red dotted positions represent unfolded assumptions and assertions, respectively. Solid and dashed arrows indicate accesses by streams and annotations, respectively



Fig. 4 A stream execution of length N + 1 with the corresponding template unfolding is depicted. The stream execution covers all cases where future out-of-accesses occur. Gray-colored and red dotted positions represent unfolded assumptions and assertions, respectively. Solid and dashed arrows indicate accesses by streams and annotations, respectively

with two dotted positions, i.e., positions where the assertion must be proven. The position where arrows are given represents the case where only the assumption results in a future out-of-bounds access. The last position of the stream execution represents the case in which both the assumption and the stream result in future out-of-bounds accesses. The presented unfolding is formalized as proof obligations in Definition 5.

Definition 5 (Proof Obligations ϕ_{End} for Future Out-ofbounds Accesses)

The proof obligations ϕ_{End} for future out-of-bounds accesses are defined as $\phi_t(p_asm, p_asserted, p_streams, p_assert)$ with template parameters:

•
$$p_asm$$
 := { $i \mid w_p \le i \le N$ },

• *p_asserted* :=
$$\{i \mid 2 \cdot w_p \le i < 3 \cdot w_p\},\$$

• p streams :=
$$\{i \mid 2 \cdot w_n \le i \le N\},$$

•
$$p_assert$$
 := { $i \mid 3 \cdot w_p \le i \le N$ }

where $N = 3 \cdot w_p + w_f$.

So far, we have defined proof obligations for certain positions in the stream execution with and without out-of-bounds accesses. Together, the proof obligations constitute an inductive argument for the correctness of the assertions, see Proposition 2. Here, the base case is given by Definition 3 and induction steps are given by Definitions 4 and 5. The induction steps use the induction hypothesis, i.e., valid assertions, due to the template parameter *p_asserted*.

Proposition 2 (Assertion Verification by LOLA Unfolding)

The set of assertions is correct if the formula $\phi_{Begin} \wedge \phi_{Run} \wedge$ ϕ_{End} is valid.

To prove that Proposition 2 holds, we distinguish exhaustively four specification cases: no temporal accesses (Proposition 3), past temporal accesses only (Proposition 4), future temporal accesses only (Proposition 5), and past and future temporal accesses (Proposition 6). First, in the case without temporal accesses, we show that all the necessary obligations for a Hoare triple are encoded in the formula and that this formula only evaluates to *false* if the assertions are not satisfied while all assumptions are.

Proposition 3 (Assertion Verification for Zero Offsets)

For a LOLA specification with $w_p = 0$ and $w_f = 0$, the set of assertions is correct if the formula $\phi_{Begin} \wedge \phi_{Run} \wedge \phi_{End}$ is valid.

Proof

The set of assertions ψ_1, \ldots, ψ_n is correct for all input streams if and only if whenever an assumption is satisfied, its corresponding assertion is satisfied as well. Without loss of generality, let $\Gamma = \{\alpha\}$, $assume(\alpha) = \{\theta\}$, and $assert(\alpha) = \{\psi\}$. We prove the proposition by showing that the formula encodes an argument for the correctness of the assertions. We consider the case that $w_p = 0$, $w_f = 0$, and the formula $\phi_{Begin} \wedge \phi_{Run} \wedge \phi_{End}$ is valid.

In this case, we do not have past or future out-of-bounds accesses. Hence, the obligations consider a single Run step with N = 0 and template parameters:

•
$$p_asm$$
 := { $i \mid i = 0$ },
• $p_asserted$:= \emptyset

•
$$p_streams$$
 := { $i \mid i = 0$ },

 $:= \{i \mid i = 0\},\$ $:= \{i \mid i = 0\}.$ • p assert

By instantiating the template for Run, we encode the following obligations:

$$\phi_{Run} \stackrel{\text{Def. 3}}{=} \phi_t \left(\begin{array}{c} p_asm, p_asserted, \\ p_streams, p_assert \end{array} \right)$$

$$\stackrel{\text{Def. 2}}{=} \bigwedge_{i \in \{0\}} smt(\theta)(i) \land \bigwedge_{i \in \emptyset} smt(\psi)(i)$$

$$\land \bigwedge_{i \in \{0\}} \left(\bigwedge_{0 < k \le n} \sigma_k = smt(e_k)(i) \right)$$

$$\rightarrow \bigwedge_{i \in \{0\}} smt(\psi)(i)$$

$$= smt(\theta)(0) \land \bigwedge_{0 < k \le n} \sigma_k = smt(e_k)(0)$$

$$\rightarrow smt(\psi)(0).$$

As can be seen, the formula encodes "assume \land program \rightarrow assertion" for a single position. Note that a single position suffices since no temporal dependencies exist. Table 1 shows that the formula evaluates to *false* only if an assertion evaluates to false, although the assumptions are valid and the output computations behave as expected. Conversely, if the

 Table 1
 Truth table for the encoding of the Hoare triple

θ	σ	ψ	$\theta \wedge \sigma \to \psi$	
0	0	0	1	assumption invalid
0	0	1	1	assumption invalid
0	1	0	1	assumption invalid
0	1	1	1	assumption invalid
1	0	0	1	outputs invalid
1	0	1	1	outputs invalid
1	1	0	0	incorrect
1	1	1	1	correct

formula is valid, then the set of assertions is correct or assumptions violated.

Note that no further obligations are added by ϕ_{Begin} and ϕ_{End} since the same template instances are created:

$$\begin{split} \phi_{Begin} &\stackrel{\text{Def. 3}}{=} \bigwedge_{0 \le N < 1} \phi_t \left(\{ i \mid 0 \le i \le N \}, \ \emptyset, \\ \{ i \mid 0 \le i \le N \}, \ \{ i \mid 0 \le i < 1 \} \right) \\ &= \phi_t (\{0\}, false, \{0\}, \{0\}) = \phi_{Run}, \\ \phi_{End} &\stackrel{\text{Def. 5}}{=} \phi_t \left(\{ i \mid 0 \le i \le 0 \}, \ \emptyset, \\ \{ i \mid 0 \le i \le 0 \}, \ \{ i \mid 0 \le i \le 0 \} \right) \\ &= \phi_t (\{0\}, false, \{0\}, \{0\}) = \phi_{Run}. \end{split}$$

Next, we show that all necessary obligations for temporal accesses to previous stream values are encoded. Further, we show that an encoding of an inductive argument is provided that considers all possible combinations of out-of-bounds accesses at the beginning of stream execution as the base case and a monitoring step with no out-of-bounds accesses as the inductive step.

Proposition 4 (Assertion Verification for Past Offsets Only)

For a LOLA specification with $w_p > 0$ and $w_f = 0$, the set of assertions is correct if the formula $\phi_{Begin} \wedge \phi_{Run} \wedge \phi_{End}$ is valid.

Example

Consider the LOLA specification

```
assume<al> reset[-1, false] \lor reset1input reset : Bool2output ol := if reset then 0 else ol[-1, 0] + 13output o2 := ol[-1, 0] + ol4assert<al> 0 \le o2 and o2 \le 35
```

that simplifies Listing 2, e.g., only past offset accesses are used. Here, w_p is 1 and w_f is 0. The unfolding of ϕ_{Begin} checks all possible combinations of out-of-bounds accesses of annotations, i.e., reset[-1, false], and outputs, i.e., o1[-1, \emptyset]. In comparison to Fig. 2, ϕ_{Begin} would only produce N = 0 and N = 1 without future accesses. Also ϕ_{Run}

represents the induction step where no default values are taken. $\hfill \Box$

Proof

The set of assertions ψ_1, \ldots, ψ_n is correct for all input streams if and only if whenever an assumption is satisfied, its corresponding assertion is satisfied as well. Without loss of generality, let $\Gamma = \{\alpha\}$, $assume(\alpha) = \{\theta\}$, and $assert(\alpha) = \{\psi\}$. We prove the proposition by showing that the formula encodes an inductive argument for the correctness of the assertions. We consider the case that $w_p > 0$, $w_f = 0$, and the formula $\phi_{Begin} \land \phi_{Run} \land \phi_{End}$ is valid.

In this case, we consider only accesses to the past. The formula encodes a *k*-induction where ϕ_{Begin} encodes the base cases and ϕ_{Run} the step case.

The template parameter for ϕ_{Begin} are:

- *p* asm := $\{i \mid 0 \le i \le N\},\$
- $p_asserted := \emptyset$,
- *p_streams* := $\{i \mid 0 \le i \le N\},$
- *p_assert* := $\{i \mid 0 \le i < \min(N+1, 2 \cdot w_p)\}$.

We have

$$\phi_{Begin} \stackrel{\text{Def. 3}}{=} \bigwedge_{0 \le N < 2 \cdot w_p} \phi_t \left(\sum_{p_streams,p_asserted, p_streams,p_assert} \right)$$

$$\stackrel{\text{Def. 2}}{=} \bigwedge_{0 \le N < 2 \cdot w_p} \left(\bigwedge_{i \in p_asm} smt(\theta)(i) \right)$$

$$\wedge \bigwedge_{i \in p_streams} \left(\bigwedge_{0 < k \le n} \sigma_k = smt(e_k)(i) \right)$$

$$\rightarrow \bigwedge_{i \in p_assert} smt(\psi)(i) \right),$$

where ϕ_{Begin} encodes stream executions that handles all possible out-of-bounds accesses. The range of stream execution N from 0 to $2 \cdot w_p$ covers all combinations of stream and assumption out-of-bounds scenarios. Two times the w_p is required since it is required in case that an output accesses available past values while at the same position the access of an assumption is out-of-bounds. Figure 5 depicts an example for $w_p = 1$. The subfigure on the top shows a stream execution where both the output computation and the assumption have out-of-bounds accesses. The subfigure on the bottom shows a stream execution where the output access to the past at the last position is in-bound but at the accessed position the access of the used assumption is not.

Furthermore, ϕ_{Run} unfolds a stream execution and uses the inductive argument. The unfolding is depicted in Fig. 6. It shows that a stream execution is unfolded with an assertion at position $3 \cdot w_p$. The inductive argument is depicted as *asserted*. The *asserted* unfolding is based on the base-cases of the *k*-induction where $k = w_p$. Further, our encoding

Both out-of-bounds:

Fig. 5 All stream executions of ϕ_{Begin} for $w_p = 1$ are depicted. It shows the case where possibly both accesses are out-of-bounds (upper) and the case where only an assumption is out-of-bounds (lower). A gray-colored box indicates that an assumption has been unfolded at this position, while a red dotted box indicates that an assertion has been unfolded at this position. Solid and dashed arrows indicate accesses by streams and annotations, respectively



Fig. 6 The unfolding of a stream execution is depicted. The assertion is proven at position $3 \cdot w_p$. Assumptions and outputs are unfolded such that all accesses can be resolved. The inductive argument that is shown in the induction base is represented by asserted

ensures that no out-of-bounds accesses occur by sufficiently unfolding the *assumptions* and *outputs*.

The template parameter for ϕ_{Run} are:

- p_asm := { $i \mid w_p \le i \le 3 \cdot w_p$ },
- *p_asserted* := $\{i \mid 2 \cdot w_p \le i < 3 \cdot w_p\},\$
- *p_streams* := $\{i \mid 2 \cdot w_p \le i \le 3 \cdot w_p\},$
- p_assert := { $i \mid i = 3 \cdot w_p$ },

and $N = 3 \cdot w_p$.

Now we obtain

$$\phi_{Run} \stackrel{\text{Def. 4}}{=} \phi_t \left(\stackrel{p_asm, p_asserted}{p_streams, p_assert} \right)$$

$$\stackrel{\text{Def. 2}}{=} \bigwedge_{i \in p_asm} smt(\theta)(i) \land \bigwedge_{i \in p_asserted} smt(\psi)(i)$$

$$\land \bigwedge_{i \in p_streams} \left(\bigwedge_{0 < k \le n} \sigma_k = smt(e_k)(i) \right)$$

$$\rightarrow \bigwedge_{i \in p_assert} smt(\psi)(i).$$

Since $w_f = 0$, no future out-of-bounds access can occur and, therefore, no obligations are added by ϕ_{End} ,

$$\phi_{End} \stackrel{\text{Def. 5}}{=} \phi_t \left(\begin{cases} i | w_p \le i \le 3 \cdot w_p \}, & \{i | 2 \cdot w_p \le i < 3 \cdot w_p \}, \\ \{i | 2 \cdot w_p \le i \le 3 \cdot w_p \}, & \{i | 3 \cdot w_p \le i \le 3 \cdot w_p \} \end{cases} \right)$$
$$= \phi_t (p_asm, p_asserted, p_streams, p_assert)$$
$$= \phi_{Run}.$$

The formula $\phi_{Begin} \wedge \phi_{Run} \wedge \phi_{End}$ intuitively encodes a *k*-induction where $k = w_p$:

$$\phi_{Begin} \begin{cases} assume(0) \land program(0) \rightarrow assert(0), \\ \dots \\ assume(k) \land program(k) \rightarrow assert(k) \end{cases}$$
$$\land \phi_{Run} \{assume(n) \land program(n) \rightarrow assert(n). \end{cases}$$

Similar to the first case, for each base case and step case, if the formula is valid then the assertions must be correct. \Box

After specifications with past offsets only, we now consider future-only specifications. Similar to before, we show that an encoding of an inductive argument is provided. In contrast to before, the encoding must also prove that default values do not violate assertions at the end of a stopped stream execution.

Proposition 5 (Assertion Verification for Future Offsets Only)

For a LOLA specification with $w_p = 0$ and $w_f > 0$, the set of assertions is correct if the formula $\phi_{Begin} \wedge \phi_{Run} \wedge \phi_{End}$ is valid.

Example

Consider the LOLA specification

<mark>assume</mark> <a1> reset ∨ reset[1, false]</a1>	1
<pre>input reset : Bool</pre>	2
<pre>output o1 := if reset then 0 else o1[1, 0] + 1</pre>	3
output $o2 := o1[1, 0] + o1$	4
assert <a1> $0 \le o2$ and $o2 \le 3$</a1>	5

that simplifies Listing 2, e.g., only future offset accesses are used. Here, w_p is 0 and w_f is 1. The unfolding of ϕ_{Begin} checks all possible combinations of out-of-bounds accesses of annotations, i.e., reset[1, false], and outputs, i.e., o1[1, 0]. ϕ_{Run} represents the induction step where no default values are taken.

Proof

The set of assertions ψ_1, \ldots, ψ_n is correct for all input streams if and only if whenever an assumption is satisfied, its corresponding assertion is satisfied as well. Without loss of generality, let $\Gamma = \{\alpha\}$, $assume(\alpha) = \{\theta\}$, and $assert(\alpha) = \{\psi\}$. We prove the proposition by showing that the formula encodes an inductive argument for the correctness of the assertions. We consider the case that $w_p = 0$, $w_f > 0$, and the formula $\phi_{Begin} \land \phi_{Run} \land \phi_{End}$ is valid.

In this case, only future stream accesses are considered. The formula encodes a *k*-induction where ϕ_{End} represent the base cases and ϕ_{Run} the step case. By unfolding a stream execution of length w_f , ϕ_{End} covers all possible future outof-bounds combinations for stream as well as annotation accesses. Figure 7 depicts ϕ_{End} where $N = w_f$. The cases are



Fig. 7 A stream executions of length w_f is depicted. It shows the case where only the last assumption access is out-of-bounds. A gray-colored box indicates that an assumption has been unfolded at this position, while a red dotted box indicates that an assertion has been unfolded at this position. Solid and dashed arrows indicate accesses by streams and annotations, respectively

similar to the base cases in the second case but this time for the future accesses. The template parameter for ϕ_{End} are:

• p_asm := { $i \mid 0 \le i \le w_f$ },

- *p_streams* := $\{i \mid 0 \le i \le w_f\},\$
- p_assert := { $i \mid 0 \le i \le w_f$ },

and $N = w_f$.

We have

$$\phi_{End} \stackrel{\text{Def. 5}}{=} \phi_t(p_asm, p_asserted, p_streams, p_assert)$$

$$\stackrel{\text{Def. 2}}{=} \bigwedge_{i \in p_asm} smt(\theta)(i) \land \bigwedge_{i \in \emptyset} smt(\psi)(i)$$

$$\land \bigwedge_{i \in p_streams} \left(\bigwedge_{0 < k \le n} \sigma_k = smt(e_k)(i) \right)$$

The induction step ϕ_{Run} unfolds a stream execution and uses the inductive argument. The unfolding is depicted in Fig. 8. In this case, the *assertion* is at the first position and the *k* base cases hold for possible future out-of-bounds accesses, i.e., at the next w_f positions. Further, *outputs* and *assumptions* are unfolded to consider all necessary *output* computations and *assertion* accesses.

The template parameter for ϕ_{Run} are:

- p_asm := { $i \mid 0 \le i \le 2 \cdot w_f$ },
- *p_asserted* := $\{i \mid 0 \le i \le w_f \land i \ne 0\},\$

 $\rightarrow \bigwedge_{i \in p_assert} smt(\psi)(i).$

- $p_streams$:= { $i \mid 0 \le i \le w_f$ },
- p_assert := { $i \mid i = 0$ },

and $N = 3 \cdot w_f$.

We thus obtain

 $\phi_{Run} \stackrel{\text{Def. 4}}{=} \phi_t(p_asm, p_asserted, p_streams, p_assert)$ $\stackrel{\text{Def. 2}}{=} \bigwedge_{i \in p_asm} smt(\theta)(i) \land \bigwedge_{i \in p_asserted} smt(\psi)(i)$ $\land \bigwedge \left(\bigwedge \sigma_k = smt(e_k)(i) \right)$

$$\rightarrow \bigwedge_{i \in p_assert}^{i \in p_streams} smt(\psi)(i).$$



Fig. 8 The unfolding of a stream execution is depicted. The assertion is proven at the first position. *Assumptions* and *outputs* are unfolded such that all accesses can be resolved. The inductive argument that is shown in the induction base is represented by *asserted*

Note that ϕ_{Begin} adds no further obligations to the formula since a single execution suffices to include all possible out-of-bounds values.

The formula $\phi_{Begin} \wedge \phi_{Run} \wedge \phi_{End}$ intuitively encodes a *k*-induction where $k = w_f$:

$$\phi_{End} \begin{cases} assume(0) \land program(0) \rightarrow assert(0), \\ \dots \\ assume(k) \land program(k) \rightarrow assert(k) \end{cases}$$

Similar to the past-only case, for each base case and step case, if the formula is valid then the assertions must be correct. \Box

Finally, the next case provides an inductive argument for LOLA specification with past and future temporal accesses. The inductive argument incorporates default values in the base case at the beginning and at the end of stream executions. Further, it handles all combinations of past and future out-of-bounds accesses.

Proposition 6 (Assertion Verification for Nonzero Offsets)

For a LOLA specification with $w_p > 0$ and $w_f > 0$, the set of assertions is correct if the formula $\phi_{Begin} \wedge \phi_{Run} \wedge \phi_{End}$ is valid.

Example

Λ

Consider the LOLA specification in Listing 2, i.e., $w_p = 1$ and $w_f = 1$. The unfolding of ϕ_{Begin} checks all possible combinations of out-of-bounds accesses of annotations, i.e., reset[-1, false] and reset[1, false], and outputs, i.e., o1[-1, 0] and o1[1, 0]. Here ϕ_{Run} represents the induction step where no default values are taken.

Proof

The set of assertions ψ_1, \ldots, ψ_n is correct for all input streams if and only if whenever an assumption is satisfied, its corresponding assertion is satisfied as well. Without loss of generality, let $\Gamma = \{\alpha\}$, $assume(\alpha) = \{\theta\}$, and $assert(\alpha) = \{\psi\}$. We prove the proposition by showing that the formula encodes an inductive argument for the correctness of the as-



Fig. 9 All combinations of past out-of-bounds accesses with future accesses are depicted

sertions. We consider the case that $w_p > 0$, $w_f > 0$, and the formula $\phi_{Begin} \land \phi_{Run} \land \phi_{End}$ is valid.

Further, we consider combinations of past and future accesses. Hence, all combinations of out-of-bounds accesses need to be covered. The formula ϕ_{Begin} covers all possible combinations of past out-of-bounds accesses with future accesses. The combinations are depicted in Fig. 9: N = 0 represents the case that all accesses are out-of-bounds, $N = w_p$ represents the case that, except for the past output stream access at the last position, all other accesses are out-of-bounds, and vice versa for the future stream access at the first position, $N = 2 \cdot w_p$ represents the case that only accesses of annotations are out-of-bounds, $N = 2 \cdot w_p + w_f$ represents both cases that only one annotation is out-of-bounds, and $N = 2 \cdot (w_p + w_f) - 1$ represents the case that only the past access of an annotation is out-of-bounds.

The template parameter for ϕ_{Begin} are:

• <i>p_asm</i>	$:= \{ i \mid 0 \le i \le N \},\$
• p_asserted	$:= \emptyset,$
• <i>p_streams</i>	$:= \{i \mid 0 \le i \le N\},\$
• p assert	$:= \{ i \mid 0 \le i < \min(N + 1, 2 \cdot w_p) \}.$

We have

$$\phi_{Begin} \stackrel{\text{Def. 3}}{=} \bigwedge_{0 \le N < 2 \cdot (w_p + w_f)} \phi_t \left(\sum_{\substack{p = asm, p = assert, \\ p = streams, p = assert}} \right)$$

$$\stackrel{\text{Def. 2}}{=} \bigwedge_{0 \le N < 2 \cdot (w_p + w_f)} \left(\bigwedge_{i \in p = asm} smt(\theta)(i) \right)$$

$$\wedge \bigwedge_{i \in p = streams} smt(\psi)(i)$$

$$\wedge \bigwedge_{i \in p = assert} smt(\psi)(i) \right).$$

Since ϕ_{Begin} covers only combinations of past out-of-bounds accesses with future accesses, ϕ_{End} covers future out-of-bounds accesses only. The unfolding is depicted in Fig. 10.

The *assertions* are unfolded such that all combinations of future out-of-bounds accesses are covered: annotation only and stream and annotation combined. Since the first *k* base cases are already covered, they can be used here as *asserted*. Further, *outputs* and *assumptions* are unfolded such that all required accesses are available.

The template parameter for ϕ_{End} are:

• p_asm	$:= \{i \mid w_p \le i \le 3 \cdot w_p + w_f\},\$
• p_asserted	$:= \{i \mid 2 \cdot w_p \le i < 3 \cdot w_p\},\$
• p_streams	$:= \{i \mid 2 \cdot w_p \le i \le 3 \cdot w_p + w_f\}$

• *p_assert* := $\{i \mid 3 \cdot w_p \le i \le 3 \cdot w_p + w_f\},\$

and $N = 3 \cdot w_p + w_f$.

We now have

 $\phi_{End} \stackrel{\text{Def. 4}}{=} \phi_t(p_asm, p_asserted, p_streams, p_assert)$

$$\stackrel{\text{Def. 2}}{=} \bigwedge_{i \in p_asm} smt(\theta)(i) \land \bigwedge_{i \in p_asserted} smt(\psi)(i)$$
$$\land \bigwedge_{i \in p_streams} \left(\bigwedge_{0 < k \le n} \sigma_k = smt(e_k)(i) \right)$$
$$\rightarrow \bigwedge_{i \in p_assert} smt(\psi)(i).$$

Last, we need to show that the assertions also hold for the case that no out-of-bounds access exists. Hence, the unfolding is encoded as depicted in Fig. 11. The *assertion* is proven at position $3 \cdot w_p$ and already shown assertions of past and future accesses are incorporated by *asserted*. Further, all required accesses of *outputs* and *assumptions* are unfolded such that no out-of-bounds access exists.

The template parameter for ϕ_{Run} are:

• p_asm := { $i \mid w_p \le i \le 3 \cdot w_p + 2 \cdot w_f$ }, • $p_asserted$:= { $i \mid 2 \cdot w_p \le i \le 3 \cdot w_p + w_f$ $\land i \ne 3 \cdot w_p$ }, • $p_streams$:= { $i \mid 2 \cdot w_p \le i \le 3 \cdot w_p + w_f$ }, • p_assert := { $i \mid i = 3 \cdot w_p$ },

and $N = 3 \cdot (w_p + w_f)$.



Fig. 10 The unfolding at the end of a stream execution is depicted. *Assumptions* and *outputs* are unfolded such that all accesses can be resolved. The inductive argument that is shown in the induction base is represented by *asserted*



Fig. 11 The unfolding of a stream execution during run is depicted. The assertion is proven at position $3 \cdot w_p$. Assumptions and outputs are unfolded such that all accesses can be resolved. The inductive argument that is shown in the induction base is represented by asserted

So we obtain

$$\phi_{Run} \stackrel{\text{Def. 4}}{=} \phi_t(p_asm, p_asserted, p_streams, p_assert)$$

$$\stackrel{\text{Def. 2}}{=} \bigwedge_{i \in p_asm} smt(\theta)(i) \land \bigwedge_{i \in p_asserted} smt(\psi)(i)$$

$$\land \bigwedge_{i \in p_streams} \left(\bigwedge_{0 < k \le n} \sigma_k = smt(e_k)(i) \right)$$

$$\rightarrow \bigwedge_{i \in p_assert} smt(\psi)(i).$$

Similar to the other cases, for each base case and step case, if the formula is valid then the assertions must be correct.

 $\begin{array}{c}
 1 \\
 2 \\
 3 \\
 4 \\
 5
 \end{array}$

By proving all cases of temporal accesses in Propositions 3 to 6, Proposition 2 is proven – the verification approach is sound. Soundness refers to the ability of an analyzer to prove the absence of errors — if a LoLA specification is accepted, it is guaranteed that the assertions are not violated. The converse does not hold, i.e., the presented verification approach is not complete. Completeness refers to the ability of an analyzer to prove the presence of errors — if a LoLA specification is rejected, the counterexample given should be a valid stream execution that results in an assertion violation. The following LOLA specification is rejected even though no assertion is violated:

Here, since the if-condition in Line 3 evaluates to *true* at the beginning of the stream execution, sum is a constant stream with value zero. Hence, the assertion in Line 4 is never violated. The verification approach rejects this specification. The reason for this is that sum ≤ 100 is added as an *asserted* condition in ϕ_{Run} . Therefore, the SMT solver can assign a value between 91 and 100 to the earliest sum variable of the unfolding, resulting in an assertion violation of the next sum variable.

4 Application experience in avionics

In this section, we present details about the tool implementation and tool experiences on practical avionic specifications.

Tool implementation and usage The tool is based on the open source RTLOLA framework.⁴ Specifically, it uses the LOLA frontend to parse a given specification into an intermediate representation. Based on this representation, the SMT formulas are created and evaluated with the Rust z3 crate.⁵ At its current phase of the crate's development, a combined solver is implemented that internally uses either a nonincremental or an incremental solver. There is no information on the implemented tactics available, but all our requests could be solved within seconds. For functions that are not natively supported by the Rust Z3 solver, the output is arbitrarily chosen by the solver with respect to the range

⁴ https://rtlola.org/. The extension is not open source yet, but will be integrated into [23].

⁵ https://docs.rs/z3/0.9.0/z3/

of the function. The tool expects a LOLA specification augmented by assumptions and assertions. The verification is done automatically and produces a counterexample stream execution, if any exists. The counterexample can then be used by the user to debug its specifications. Two different kinds of user are targeted. First, users that write the entire augmented specification. Such a user could be a system engineer who is developing a safety monitor and wants to ensure that it contains critical properties. Second, users that augment an existing specification. Here, one reason could be that an existing monitor shall be composed with other critical components and certain behavioral properties are expected. Also, similar to software testing, the task of writing a specification and their respective assumptions and assertions could be separated between two users to ensure the independence of both.

Practical results To gain practical tool experience, previously written specifications based on interviews with engineers of the German Aerospace Center [25] were extended by assumptions and assertions. The previous specifications were tested using log-files and simulations - the authors considered them correct. We report several specification errors in Table 2 that were detected by the presented verification extension. In fact, the detected errors would have resulted in undetected failures. After the errors in the previous specifications were fixed, all assertions were proven correct. Note that the errors could have been found by manual reviews. However, such reviews are tedious and error-prone, especially when temporal behaviors are involved. The detected errors in Table 2 can be grouped into three classes: Classical Bugs, Operator Errors, and Wrong Interpretations. Classical bugs are errors that occur when implementing an algorithm. Operator errors are LOLA specific errors, e.g., temporal accesses. Last, wrong interpretations refer to gaps between the specification and the user's design intend, e.g., violated assertions due to incomplete specifications. Next, we give one representative example for each group. We reduced the specification to the representative fragment.

Example 1 (Classical Bug)

The LOLA specification in Listing 3 monitors the fuel level. A monitor shall notify the operator when one of the three different fuel levels are reached: half (Line 9), warning (Line 10), and danger (Line 11). The fuel level is computed as a percentage in Lines 7 to 8. It uses the fuel level at the beginning of the flight (Line 6) as a reference for its computation. Given the documentation of the fuel sensor, it is known that fuel values are within \mathbb{R}^+ and decreasing. This is formalized in Line 4 as an assumption. As an invariant, we asserted that the starting fuel is greater or equal to fuel (Line 16). Further, in Lines 17 to 19, we stated that once a

// Inputs	1
input fuel: Float64	2
// Assumptions	3
assume <a5> fuel > 0.0 and fuel < fuel[-1, fuel + 0.1]</a5>	4
// Outputs	5
<pre>output start_fuel := start_fuel[-1, fuel]</pre>	6
<pre>output fuel_level :=</pre>	7
<pre>1.0 - (start_fuel - fuel) / start_fuel</pre>	8
<pre>output fuel_half := fuel_level < 0.50</pre>	9
<pre>output fuel_warning := fuel_level < 0.25</pre>	10
<pre>output fuel_danger := fuel_level < 0.10</pre>	11
<pre>trigger_once fuel_half "INFO: Fuel is below 50%"</pre>	12
<pre>trigger_once fuel_warning "WARNING: Fuel is below 25%"</pre>	13
<pre>trigger_once fuel_danger "DANGER: Fuel is below 10%"</pre>	14
// Assertions	15
<pre>assert<a5> start_fuel >= fuel</a5></pre>	16
and (fuel_half[-1, false] -> fuel_half)	17
and (fuel_warning[-1, false] -> fuel_warning)	18
and (fuel danger[-1, false] -> fuel danger)	19

Listing 3 The fixed version of the LOLA ctrl_output specification that monitors the fuel level. Three level of engagement are depicted: half, warning, and danger

Table 2 Detected errors by the verification extension, where #o, #a, and #g represent the number of outputs, assumptions, and assertions given in the specification, respectively

Specification	Appx.	#o	#a	#g	Detected errors	
gps_vel_output	A.1	14	6	6	_	
gps_pos_output	A.2	19	3	10	_	
imu_output	A.3	18	6	6	Wrong default value Division by zero	
nav_output	A.4	25	3	5	Missing abs()	
tagging	A.5	6	2	2	_	
ctrl_output	A.6	25	7	8	Wrong threshold comparisons	
mm_output_1	A.7	4	1	2	-	
mm_output_2	A.8	17	6	9	Missing if condition Wrong default value	
contingency_output	A.9	4	8	1	Observation: both contingencies could be true in case of voting, i.e., both at 50%	
health_output	A.10	1	5	1	_	

level is reached it should remain at this level. During our experiment, the assertion led to a counterexample that pointed to the previously used and erroneous fuel level computation:

```
output fuel_level := (start_fuel - fuel) / start_fuel
```

In short, the output computed the consumed fuel and not the remaining fuel. The computation could be easily fixed by converting consumed fuel into remaining fuel, see Line 8. Therefore, Listing 3 satisfies its assertion. Note, that offset accesses were used to assert the temporal behavior of the fuel level output stream. Further, trigger_once is an abbreviation which states that only the first raising edge is reported to the user.

Example 2 (Operator Error)

An important monitoring property is to detect frozen values as these indicate a deteriorated sensor. Such a specification is depicted in Listing 4. Here, as an input, the acceleration in xdirection is given. The frozen value check is computed from Line 6 to 10. It compares previous values using LoLA's offset operator. To check this computation, we added the sanity check that asserts that no frozen value shall be detected (Line 13) when small changes in the input are present (Line 4). In the previous version, the frozen values were computed using the abbreviated offset operator:

output frozen_ax := ax[-5..0, 0.0, =]

This resulted in a counterexample that pointed to wrong default values. Although the abbreviated version is easier to read and reduces the size of the specification, it is unfortunately not suitable for this kind of property. The tool detected the unlikely situation that the first value of ax is 0.0 which would have resulted in evaluating frozen_ax to true. Although unlikely, this should be avoided as contingencies activated in such situations depend on correct results and otherwise could harm people on the ground. By unfolding the operator and adding a different default value to one of the past accesses, the error was resolved (Line 6). Listing 4 shows the fixed version which satisfies its assertion.

```
// Inputs
input ax: Float32
                                                             2
                                                             3
// Assumptions
assume <a1> ax != ax[-1, ax + \varepsilon]
                                                             4
                                                             5
// Outputs
output frozen_ax := ax[-5, 0.1] = ax[-4, 0.0]
                                                             6
                 and ax[-4, 0.0] = ax[-3, 0.0]
                                                             7
                 and ax[-3, 0.0] = ax[-2, 0.0]
                                                             8
                 and ax[-2, 0.0] = ax[-1, 0.0]
                                                             0
                 and ax[-1, 0.0] = ax
trigger frozen_ax "WARNING: x-acceleration is frozen!"
                                                             11
// Assertions
                                                              12
assert <a1> !frozen_ax
                                                             13
```

Listing 4 The LOLA imu_output specification that monitors frozen acceleration values

Example 3 (Wrong Interpretation)

In Listing 5, two visual sensor readings are received (Lines 2-5). Both, readings argue over the same observations where avgDist represents the average distance to the measured obstacle, actual is the number of measurements, and static is the number of unchanged measurements. A simple rating function is introduced (Lines 7-12) that estimates the corresponding rating – the higher the better. Using these ratings, the trust in each of the sensors is computed probabilistically (Lines 11-13). When considering the integration of such a monitor as an ASTM F3269-21 switch condition that decides which sensor value should be forwarded, the specification should be revised. This is the case because the assertion in Line 17 produces a counterexample which indicates that both trust triggers (Lines 14-15) can be activated at the same time. A common solution for this problem is to introduce a priority between the sensors.

The examples show how the presented LOLA verification extension can be used to find errors in specifications. We also noticed that the annotations can serve as documentation. System assumptions are often implicitly known during development and are finally documented in natural language in separate files. Having these assumptions explicitly stated within the monitor specification potentially reduces future mistakes when reusing the specification, e.g., when composing with other monitor specifications. Listing 6 depicts such an example specification. Here, the monitor interfaces are clearly defined by the domain of input *a* (Line 5) and output *o* (Line 13). Also, *reset* is assumed to be valid at least once per second (Line 5). Further, no deeper understanding of the internal computations (Lines 7–10) is required in order to safely compose this specification with others.

5 Monitoring assumptions

The presented verification approach offers an analysis of the specification that can guarantee the desired behavior of the

// Inputs	1
<pre>input avgDist_laser, actual_laser,</pre>	2
<pre>static_laser: Float64</pre>	3
<pre>input avgDist_optical, actual_optical,</pre>	4
static_optical: Float64	5
// Outputs	6
<pre>output rating_laser := 0.2 * static_laser</pre>	7
+ 0.4 * actual_laser + 0.4 * avgDist_laser	8
<pre>output rating_optical := 0.2 * static_optical</pre>	9
+ 0.4 * actual_optical + 0.4 * avgDist_optical	10
<pre>output trust_laser :=</pre>	11
<pre>rating_laser / (rating_laser + rating_optical)</pre>	12
<pre>output trust_optical := 1.0 - trust_laser</pre>	13
<pre>trigger trust_laser >= 0.5</pre>	14
<pre>trigger trust_optical >= 0.5</pre>	15
// Assertions	16
<pre>assert <a1> trust_laser != trust_optical</a1></pre>	17

Listing 5 The LOLA contingency_output specification that uses an heuristic to decide which sensor is more trustworthy

```
// Inputs with frequency 5Hz
input a: Float64
                                                                2
input reset: Bool
                                                               3
// Assumptions
                                                                4
assume <a1> 0.0 \leq a \leq 1.0 and reset[-4..0, false, \lor]
// Outputs
output o_1 := ...
                                                               9
output o n := \dots
output o := o_1 + \ldots + o_n
                                                                10
trigger o \geq 0.5 "Warning: Output o exceeds threshold!"
                                                                11
                                                                12
// Assertions
assert <a1> 0.0 \leq o \leq 1.0
                                                               13
```

Listing 6 LOLA specification annotations describe interface properties

monitor outputs. Yet, these guarantees are often based on assumptions that can be violated at runtime. Further, a violated assumption does not directly result in a violated assertions. As an example consider Listing 4 that states an assertion which is violated when six consecutive **ax** values are the same. Hence, the sequence $\langle 0.1; 0.1; 0.1; 0.4; 0.5 \rangle$ satisfies the assertion. Yet, it violates the assumption at the first positions.

In this section, we will give a translation of an annotated specification into a specification that checks assumptions at runtime and efficiently activates and deactivates assertion checks. Further, we will present experimental results showing that the verification extension not only provides static guarantees, but that translating it into a corresponding specification can lead to better runtime performance compared to a monitor that simply checks all assertions during runtime.

5.1 Translation into Lola 2.0

We translate an annotated LOLA specification into a LOLA 2.0 specification [10]. LOLA 2.0 extends LOLA by streams that can be spawned, filtered, and closed at runtime. For the translation, we only need to handle assumptions and assertions since trigger, input, and output streams are directly supported by LOLA 2.0.

We replace each assumption by an output stream. A stream **output** o : Bool spawn if e_s filter e_f

close $e_c := e(t_1, \ldots, t_m, s_1, \ldots, s_n)$ is an output stream that is created when its spawn condition e_s is true. It then starts producing values by evaluating its computation e if its filter condition e_f is *true* until its close condition e_c is satisfied. Since a violated assumption can influence previous and future assertions due to temporal offset accesses, the computation e is a counter that represents how many assertions are impacted by the assumption. If the counter is positive, then an assumption was violated that influences an assertion computation. To compute the impact of a violated assumption, we take the maximum between the longest chain of offset accesses from assertion to assumption plus one. This is achieved by an analysis of the dependency graph [8]. To start the counter, we use the negated assumption as spawn condition. Further, we close the stream when the impact of a violated assumption is over, i.e., when the counter is zero. As filter condition, we use *true* to decrease the counter in each step.

We also represent assertion checks by output streams. For each assertion, we use an output stream that is only extended if one of its corresponding assumptions is violated, i.e., its counter value is positive. We also add a trigger to report assertion violations.

As an example consider the annotated LOLA specification

<pre>input vel : Float32 // Velocity</pre>	1
assume <a> -20.0 \leq vel \leq 20.0	2
<pre>output vel_max := max(abs(vel), vel_max[-1, 0.0])</pre>	3
<pre>trigger vel_max > 20.0 "Velocity threshold exceeded!"</pre>	4
assert <a> abs(vel[-20, 0.0, +]) / $3.0 \leq 20.0$	5

that checks the maximal velocity value (Line 3) and the average velocity over a discrete window of three (Line 5). The assumption (Line 2) and the assertion (Line 5) are transformed to

output assumption	1
spawn if !(-20.0 \leq vel \leq 20.0)	2
<pre>filter true close assumption = 0</pre>	3
:= if -20.0 \leq vel \leq 20.0	4
then assumption[-1,0] - 1 else 1	5
<pre>trigger assumption > 0 "Assumption violated!"</pre>	6
output assertion	7
<pre>spawn if true filter assumption > 0 close false</pre>	8
:= abs(vel[-20, 0.0, +]) / $3.0 \le 20.0$	9
trigger !assertion "Assertion violated!"	-10

The output *vel_max* and trigger remain unchanged. Note that the trigger could have been replaced by an assertion. Yet, this would not reduce as much overhead as for the window check which the following experiments will show.

5.2 Experiments

For our experiments, we compare the performance of the presented translation to a naive translation that checks assumptions and assertions independently in each execution step. As annotated specification, we use

that we scale in the number of annotation pairs using the variable *i* and the computational load of the assertion by the window variable *w*. For instance, i = 10 and w = 5 produces ten inputs with the corresponding annotations where each assertion takes the sum over the last five input values including the current one. The naive translation v_n with omitted triggers is

```
input a_i: Float64 1
output assumption<sub>i</sub> := a_i \le 2.0 2
output assertion<sub>i</sub> := a_i [-w..0, 0.0, +] \le w \cdot 2.0 3
```

The presented translation v_t also with omitted triggers is

```
input a_i: Float64
output assumption;
   spawn if a_i > 2.0 filter true close assumption; = 0
:= if a_i \le 2.0 then assumption; [-1,0] - 1 else w
output assertion;
   spawn if true filter assumption; > 0 close false
:= a_i [-w..0, 0.0, +] < w · 2.0</pre>
```

For the experiments, the considered values of i were 5, 10, and 15 and the values of w were 0, 5, and 10. The experiments were conducted on three different kinds of log-files: no assumption is violated, all assumptions are violated, half of the assumptions are violated. Each log-file contains 10.000.000 events that were sufficient to report the average time in nanoseconds required by the monitor to evaluate one input event. Each experiment was carried out three times and the average was taken. For the experiments, an eight-

core machine with an 2.5 GHz Intel i7 processor with 32 GB RAM was used.

The results of the experiments are depicted in Fig. 12. As can be seen in Fig. 12a, which considers log-files with no violation of assumptions, version v_t significantly improves runtime by up to 64.06%. It can also be seen that version v_t improves the required time per event by 8.17% already in the case of simple assertions. Further, the required time for v_t remained constant while increasing the window size which shows that no unnecessary assertion checks were computed; in contrast to v_n , where the required time correlates with the size of the window. Next, Fig. 12b considers log-files where all assumptions are violated. The results show that this time v_t correlates with the size of the window similar to v_n since all the assertions need to be checked due to violated assumptions. The experiments show that v_t incurs an overhead of up

	Number of inputs i									
Window w	5			10			15			
	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	$\mathbf{\Delta}$ [%]	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	$\mathbf{\Delta}$ [%]	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	$\mathbf{\Delta}$ [%]	
0	1203.67	1105.34	8.17	2404.67	2096.67	12.81	3674.67	3236.34	11.93	
5	1951.67	1118.00	42.72	4146.34	2127.67	48.69	6318.00	3273.34	48.19	
10	2859.67	1124.34	60.68	5985.00	2151.00	64.06	8977.00	3228.00	64.04	

3

4

	Number of inputs i									
Window w		5		10			15			
	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	Δ [%]	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	$\mathbf{\Delta}$ [%]	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	Δ [%]	
0	2371.34	3092.67	-30.42	4637.67	6165.34	-32.94	6922.34	9266.34	-33.86	
5	3124.00	3881.34	-24.24	6376.00	7956.34	-24.79	9440.34	11842.34	-25.44	
10	3972.34	4844.00	-21.94	8203.67	9852.00	-20.09	12157.00	14513.67	-19.39	

(a) None of the 10.000.000 events in the log-file violates the assumption.

				Num	ber of in	puts i			
Window w	5			10			15		
	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	$\mathbf{\Delta}$ [%]	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	$\mathbf{\Delta}$ [%]	$\mathbf{v_n} \ [\mu s]$	$\mathbf{v_t} \ [\mu s]$	$\mathbf{\Delta}$ [%]
0	1783.67	2099.00	-17.68	3505.00	4163.00	-18.77	5311.00	6270.34	-18.06
5	2535.00	2515.00	0.79	5222.34	5043.00	3.43	7839.00	7518.34	4.09
10	3409.34	2967.00	12.97	7071.00	6032.34	14.69	10560.67	8880.67	15.91

(b) All of the 10.000.000 events in the log-file violate the assumption.



(c) Half of the 10.000.000 events in the log-file violate the assumption.

(d) Graphical representation of Figure 12a on the left, 12b in the middle, and 12c on the right.

Fig. 12 The results of the log-file analyses using the translations of an annotated specification is given. Entries in the table represent the time required by the monitor for one input event. The specification version v_n represents a specification that checks assumptions and assertion for each event in the log-file whereas the specification version v_t checks

an assertions only if its corresponding assumption is violated by the use of output streams that use spawn, filter, and close conditions. The symbol Δ represents the runtime effect of dynamic assertion checks, i.e., positive values indicate improvement and negative values indicate deterioration

to -33.86%. However, Fig. 12c shows that already in the case where half of the inputs violate the assumptions and a more complex assertion is used (w = 5), the LOLA 2.0 specification version v_t pays off and outperforms v_n by up to 15.91%. The results are also graphically depicted in Fig. 12d.

Overall, the experiments show that translating an annotated specification into a LOLA 2.0 specification can be used to report assertion violations due to violated assumptions efficiently at runtime. Since assumptions are generally expected to be satisfied in the nominal case, the translation also improves the monitor's runtime without losing its guarantees. Especially complex assertions based on simple assumptions benefit from the translation. If the assertions are simple, the benefits from the translation are negligible.

Remark

We also considered the alternative LOLA assumption encoding

```
output assumption1:= if a_i \leq 2.0 then if assumption_i[-1,0] = 0 then 02else assumption_i[-1,0] - 1 else w
```

that no longer uses LOLA 2.0 features. Our result showed a runtime improvement of up to 59.17% in the case of no violations and a runtime deterioration of only up to -8.54%. Still, we decided on the LOLA 2.0 assumption encoding to gain the best performance, since assumptions should not be violated in the nominal case. Yet, these results indicate that the parameterization of the assumptions has the largest share in the reported deterioration.

6 Conclusion

As both the relevance and the complexity of cyber-physical systems continue to grow, runtime monitoring is an essential ingredient of safety-critical systems. When monitors are derived from specifications it is crucial that the specifications are correct. In this paper, we have presented a sound verification approach for the stream-based monitoring language LOLA. With this approach, the developer can formally prove guarantees on the streams computed by the monitor, and hence ensure that the monitor does not cause dangerous situations. The verification extension is motivated by upcoming aviation regulations and standards as well as by practical feedback of engineers.

The extension has been applied to previously written LOLA specifications that were obtained based on interviews with aviation experts. In this process, we discovered and fixed several serious specification errors.

Further, since assumption can fail during runtime, they must be monitored and only when they are violated, their respective assertions need to be monitored as well. In this paper, we have efficiently monitored verified guarantees at runtime. Our experiments have shown that our LOLA 2.0 encoding can significantly improve the monitors performance while maintaining a low overhead in case of few assumption violations. Yet, this improvement is highly dependent on the given specification. In general, simple assumptions and complex assertions benefit from this approach.

In the future, we plan to develop automatic invariant generation for LOLA specifications. Another interesting direction for future work is to support the effort of [2] by exploiting the results of the analysis for the optimization of the specification and the resulting monitoring code. Finally, we plan to extend the verification approach to RTLOLA, the real-time extension of LOLA.

Appendix: Lola specifications – experience report

A.1 Specification: gps_vel_output

input sol_age: Float32	1
input hor_spd: Float32	2
input trk_gnd: Float32	3
input vert_spd: Float32	4
input time_s: UInt64	5
input time_us: UInt64	6
// Assumptions	7
assume <a1> time – time.offset(by: –1).defaults(to:</a1>	8
time -0.1 > 0.0	
and time – time.offset(by: –1).defaults(to: time –	9
$(0.1) \le 0.1$	
and trace_pos ≥ 0	10
	11
assume <a2> time – time.offset(by: –1).defaults(to:</a2>	12
time -0.1 > 0.0	
and time - time.offset(by: -1).defaults(to: time -	13
$(0.1) \le 0.1$	
and trace_pos ≥ 0	14
// Frequency computations	15
output time := cast(time_s) + cast(time_us) / 1000000.0	16
<pre>output start_time := if time.offset(by: -1).defaults(to:</pre>	17
-1.0 = -1.0 then time else start_time.offset(by:	
-1).defaults(to: -1.0)	
<pre>output flight_time := time - start_time</pre>	18
output trace_pos @ sol_age or hor_spd or trk_gnd or	19
vert_spd or time_s or time_us :=	
trace_pos.offset(by: -1).defaults(to: -1) + 1	
output frequency :=	20
1.0 / (time – time.offset(by: –1).defaults(to: time	21
- 0.0001))	
<pre>output freq_sum :=</pre>	22
<pre>freq_sum.offset(by: -1).defaults(to: 0.0) +</pre>	23
frequency	
<pre>output freq_avg := freq_sum / cast(trace_pos+1)</pre>	24

<pre>output freq_max := if frequency > freq_max.offset(by:</pre>	25
output freq_min := if frequency < freq_min.offset(by: -1).defaults(to: frequency) then frequency else freq_min offset(by: -1) defaults(to: frequency)	26
// Speed computations	27
output hor_spd_max := if hor_spd >	28
hor_spd_max.offset(by: -1).defaults(to: 0.0) then	
<pre>hor_spd else hor_spd_max.offset(by:</pre>	
-1).defaults(to: 0.0)	
<pre>output vert_spd_max := if vert_spd ></pre>	29
vert_spd_max.offset(by: -1).defaults(to: 0.0) then	
vert_spd_else_vert_spd_max.offset(by:	
-1).defaults(10: 0.0)	30
wrt_north)	50
trigger sol age ≤ 0.5 "Sol age should remain zero!"	31
output trk gnd in bound := if trk gnd ≥ 0.0 and	32
trk_gnd <= 360.0 then	
<pre>trk_gnd_in_bound.offset(by: -1).defaults(to:</pre>	
true) else false	
<pre>output trk_gnd_max := if trk_gnd ></pre>	33
trk_gnd_min.offset(by: -1).defaults(to: 0.0) then	
trk_gnd else trk_gnd_min.offset(by:	
-1).defaults(to: 0.0)	24
trk and max offset(by: -1) defaults(to: 0.0) then	34
trk ond else trk ond max offset(by:	
-1).defaults(to: 0.0)	
// Assertions	35
assert <a1> time.offset(by: -1).defaults(to: -1.0) <</a1>	36
time	
and start_time == start_time.offset(by:	37
-1).defaults(to: start_time)	
and flight_time >= flight_time.offset(by:	38
-1).defaults(to: 0.0)	20
assert $\langle a2 \rangle$ inequency $\rangle = 10.0$	- 39 - 40
-1) defaults(to: 0.0) + 10.0	40
assert <a3> trk gnd in bound.offset(by:</a3>	41
-1).defaults(to: true)	
or !trk_gnd_in_bound	42

A.2 Specification: gps_pos_output

import math	1
input lat: Float32	2
input lon: Float32	3
input hgt: Float32	4
input nObjs: UInt64	5
input nGPSL1: UInt64	6
input time_s: UInt64	7
input time_us: UInt64	8
// Assumptions	9
assume <a1> time – time.offset(by: –1).defaults(to:</a1>	10
time -0.1 > 0.0	
and time - time.offset(by: -1).defaults(to: time -	11
$(0.1) \le 0.1$	

and trace_pos ≥ 0	12
// Frequency computations	13
output time: Float32 := cast(time_s) + cast(time_us) / 1000000.0	14
<pre>output start_time := if time.offset(by: -1).defaults(to:</pre>	15
-1.0 == -1.0 then time else start_time.offset(by:	
-1).defaults(to: -1.0)	
<pre>output flight_time := time - start_time</pre>	16
output trace_pos @ lat or lon or hgt or nObjs or	17
nGPSL1 or time_s or time_us :=	
trace_pos.offset(by: -1).defaults(to: -1) + 1	1.0
$\frac{1}{2} defaults(to: time = 0.0001))$	18
-1).ucrauits(io. time - 0.0001))	10
-1) defaults(to: 0.0) + frequency	17
output freq avg := freq sum / cast(trace $pos+1$)	20
output freq_arg : neq_sum / cus(arcc_pos+1) output freq_max := if frequency > freq_max.offset(by:	21
-1).defaults(to: 0.0) then frequency else	
freq_max.offset(by: -1).defaults(to: 0.0)	
output freq_min := if frequency < freq_min.offset(by:	22
-1).defaults(to: 0.0) then frequency else	
<pre>freq_min.offset(by: -1).defaults(to: 0.0)</pre>	
// Statistics	23
<pre>output lat_max := if lat > lat_max.offset(by:</pre>	24
-1).defaults(to: lat) then lat else	
lat_max.offset(by: -1).defaults(to: lat)	
<pre>output lat_min := if lat < lat_min.offset(by:</pre>	25
-1).defaults(to: lat) then lat else	
lat_min.offset(by: -1).defaults(to: lat)	26
output $Ion_max := II Ion > Ion_max.offset(by: 1) defoults(ter, lon) then lon also$	20
-1).defaults(to: 1011) then 1011 else	
autnut lon min := if lon < lon min offset(by:	27
-1).defaults(to: lon) then lon else	21
lon min.offset(by: -1).defaults(to: lon)	
output lat in bound := max(abs(lat max),	28
abs(lat_min)) <= 90.0	
<pre>output lon_in_bound := max(abs(lon_max),</pre>	29
abs(lon_min)) <= 180.0	
<pre>trigger !lat_in_bound "Irregular latitude value!"</pre>	30
trigger !lon_in_bound "Irregular longitude value!"	31
output begin := false	32
<pre>output start_height := if begin.offset(by:</pre>	33
-1).defaults(to: true) then hgt else	
start_height.offset(by: -1).defaults(to: 0.0)	24
1) defaulte(to: 0.0) hetstart_height)	34
- 1).defaults(to. 0.0), ligt - start_lieight)	35
-1) defaults(to: 0.0) hot - start height)	55
trigger het inc max > 100.0 "Never increase height by	36
more than 100m!"	00
trigger hgt dec max < -100.0 "Never decrease height	37
by more than 100m"	
// Assertions	38
assert <a1> time.offset(by: -1).defaults(to: -1.0) <</a1>	39
time	
and start_time == start_time.offset(by:	40
-1).defaults(to: start_time)	
and flight_time >= flight_time.offset(by:	41
-1).defaults(to: 0.0)	

assert <a2> hgt_inc_max >= 0.0 and hgt_dec_max <=</a2>	42
0.0	
and hgt_inc_max >= hgt_inc_max.offset(by:	43
-1).defaults(to: 0.0)	
and hgt_dec_max <= hgt_inc_max.offset(by:	44
-1).defaults(to: 0.0)	
and start_height = start_height.offset(by:	45
-1).defaults(to: start_height)	
and (lat_in_bound.offset(by: -1).defaults(to: true) or	46
!lat_in_bound)	
and (lon_in_bound.offset(by: -1).defaults(to: true) or	47
!lon_in_bound)	

A.3 Specification: *imu_output*

import math	1
input ax: Float32	2
input ay: Float32	3
input az: Float32	4
input time_s: UInt64	5
input time_us: UInt64	6
input counter: Int64	7
// Assumptions	8
assume <a1> time – time.offset(by: –1).defaults(to:</a1>	9
time -0.1 > 0.0	
and time – time.offset(by: -1).defaults(to: time –	10
$(0.1) \le 0.1$	
and trace_pos ≥ 0	11
assume <a2> ax != ax.offset(by: -1).defaults(to: ax +</a2>	12
0.1)	
and ay $!=$ ay.offset(by: -1).defaults(to: ay + 0.1)	13
and az $!=$ az.offset(by: -1).defaults(to: az + 0.1)	14
// Frequency computations	15
<pre>output time := cast(time_s) + cast(time_us) / 1000000.0</pre>	16
<pre>output start_time := if time.offset(by: -1).defaults(to:</pre>	17
-1.0 == -1.0 then time else start_time.offset(by:	
-1).defaults(to: -1.0)	
<pre>output flight_time := time - start_time</pre>	18
output trace_pos @ ax or ay or az or time_s or time_us	19
or counter := trace_pos.offset(by: -1).defaults(to:	
-1) + 1	
output frequency := 1.0 / (time – time.offset(by:	20
-1).defaults(to: time -0.0001))	
<pre>output freq_sum := freq_sum.offset(by:</pre>	21
-1).defaults(to: 0.0) + frequency	
<pre>output freq_avg := freq_sum / cast(trace_pos+1)</pre>	22
// Statistics	23
output deviation := abs(frequency – 100.0)	24
<pre>output exceeds_worst := deviation ></pre>	25
<pre>worst_dev.offset(by: -1).defaults(to: 0.0)</pre>	
<pre>output worst_dev_pos := if exceeds_worst then</pre>	26
<pre>trace_pos else worst_dev_pos.offset(by:</pre>	
-1).defaults(to: 0)	
<pre>output worst_dev := if exceeds_worst then deviation</pre>	27
else worst_dev.offset(by: -1).defaults(to: 0.0)	
<pre>output ax_max :=</pre>	28
<pre>max(abs(ax),ax_max.offset(by:-1).defaults(to:0.0))</pre>	
<pre>output ay_max :=</pre>	29
<pre>max(abs(ay),ay_max.offset(by:-1).defaults(to:0.0))</pre>	

<pre>output az_max :=</pre>	30
max(abs(az),az_max.offset(by:-1).defaults(to:0.0))	
trigger ax > 15.0 or ay > 15.0 or az > 15.0	31
output frozen_ax := ax.offset(by:-1).defaults(to:0.0) =	32
ax	
and $ax.offset(by:-2).defaults(to:0.0)=$	33
ax.offset(by:-1).defaults(to:0.0)	
and $ax.offset(by:-3).defaults(to:0.0) =$	34
ax.offset(by:-2).defaults(to:0.0)	
and ax.offset(by:-4).defaults(to:0.0)=	35
ax.offset(by:-3).defaults(to:0.0)	
and ax.offset(by:-5).defaults(to:0.1)=	36
ax.offset(by:-4).defaults(to:0.0)	
output frozen_ay := ay.offset(by:-1).defaults(to:0.0) =	37
ay	
and ay.offset(by:-2).defaults(to:0.0)=	38
ay.offset(by:-1).defaults(to:0.0)	
and ay.offset(by:-3).defaults(to:0.0)=	39
ay.offset(by:-2).defaults(to:0.0)	
and ay.offset(by:-4).defaults(to:0.0)=	40
ay.offset(by:-3).defaults(to:0.0)	
and ay.offset(by:-5).defaults(to:0.1)=	41
ay.offset(by:-4).defaults(to:0.0)	
output frozen_az := az.offset(by:-1).defaults(to:0.0) =	42
az	
and az.offset(by:-2).defaults(to:0.0)=	43
az.offset(by:-1).defaults(to:0.0)	
and az.offset(by:-3).defaults(to:0.0)=	44
az.offset(by:-2).defaults(to:0.0)	
and az.offset(by:-4).defaults(to:0.0)=	45
az.offset(by:-3).defaults(to:0.0)	
and az.offset(by:-5).defaults(to:0.1)=	46
az.offset(by:-4).defaults(to:0.0)	
trigger frozen_ax or frozen_ay or frozen_az	47
output check_counter := if trace_pos = 0 then false	48
else (counter != (counter.offset(by:	
-1).defaults(to: -1) + 1) % 100)	
trigger check_counter "A counter value was ignored."	49
// Assertions	50
assert <a1> time.offset(by: -1).defaults(to: -1.0) <</a1>	51
time	
and start_time == start_time.offset(by:	52
-1).defaults(to: start_time)	
and flight_time >= flight_time.offset(by:	53
-1).defaults(to: 0.0)	
assert <a2> !frozen_ax and !frozen_ay and !frozen_az</a2>	54

A.4 Specification: *nav_output*

import math	1
input lat: Float32	2
input lon: Float32	3
input ug: Float32	4
input vg: Float32	5
input wg: Float32	6
input time_s: UInt64	7
input time_us: UInt64	8
// Assertion	9

assume $$ trace pos $>=0$	10
and time – time. $\overline{offset}(by: -1)$. $defaults(to: time –$	11
$(0.1) \le 0.1$	
and time – time.offset(by: -1).defaults(to: time –	12
(0.1) > 0.0	
// Frequency Computation	13
output time := cast(time_s) + cast(time_us) / 1000000.0	14
output start_time := if time.offset(by: -1).defaults(to:	15
-1.0 == -1.0 then time else start_time.offset(by:	
-1).defaults(to: -1.0)	
output flight_time := time - start_time	16
output trace_pos @lat or lon or ug or vg or wg or	17
time_s or time_us := trace_pos.offset(by:	
-1).defaults(to: -1) + 1	10
$\frac{1}{2} defaults(te: time = 0.0001))$	10
-1).defaults(to: time -0.0001)	10
-1) defaults(to: 0.0) + frequency	12
output freq avg := freq sum / cast(trace $pos+1$)	20
output freq_avg := if frequency > freq_max offset(by:	21
-1) defaults(to: frequency) then frequency else	- 1
freq max.offset(by: -1).defaults(to: frequency)	
output freq min := if frequency < freq min.offset(by:	22
-1).defaults(to: frequency) then frequency else	
freq_min.offset(by: -1).defaults(to: frequency)	
// Statistics	23
output velocity := sqrt(ug*ug + vg*vg + wg*wg)	24
output lon1_rad := lon.offset(by: -1).defaults(to: 0.0)	25
* 3.1415926535 / 180.0	
output lon2_rad := lon * 3.1415926535 / 180.0	26
output lat1_rad := lat.offset(by: -1).defaults(to: 0.0) *	27
3.1415926535 / 180.0	
output lat2_rad := lat * 3.1415926535 / 180.0	28
output dlon := lon2_rad - lon1_rad	29
output dlat := $lat2_rad - lat1_rad$	30
$\begin{array}{l} \text{output a} := (\sin((1at/2.0)) * (\sin((1at/2.0)) + \\ \cos((1at/2.0)) + \\ \cos$	31
$\cos(a _1 - a) \approx \cos(a _2 - a) \approx \cos(a _2 - a)$	
$(\sin(\alpha \cos 2.0)) * (\sin(\alpha \cos 2.0))$	37
output x_atan2 := $sqrt(a)$	33
output $c := 2.0 * \text{if } x \text{ atan}^2 > 0.0 \text{ then}$	34
arctan(v atan2/x atan2)	0.
else if x_atan2 < 0.0 and y_atan2 >= 0.0	35
then $\arctan(y_{atan}^{2}/x_{atan}^{2}) + 3.1415926535$	36
else if $x_{atan2} < 0.0$ and $y_{atan2} < 0.0$	37
then arctan(y_atan2/x_atan2) - 3.1415926535	38
else if $x_{atan2} = 0.0$ and $y_{atan2} > 0.0$ then	39
3.1415926535 / 2.0	
else if $x_{atan2} = 0.0$ and $y_{atan2} < 0.0$ then	40
-3.1415926535 / 2.0	
else 0.0	41
output gps_distance := 6373000.0 * c	42
output passed_time := time - time.offset(by:	43
-1).defaults(to: 0.0)	л л
output distance_max := velocity * passed_time	44
distance max)	43
usually_max j	46
dif distance>1	+0
trigger detected jump "Tump!"	47
	-r /

// Assertions	48
time	49
and start_time == start_time.offset(by:	50
-1).defaults(to: start_time)	
and flight_time >= flight_time.offset(by:	51
-1).defaults(to: 0.0)	
<pre>assert <a2> (!detected_jump or gps_distance ></a2></pre>	52
distance_max)	
or (!detected_jump or distance_max > gps_distance)	53

A.5 Specification: *tagging*

import math	1
input time_s: UInt64	2
input time_us: UInt64	3
input vel: Float64	4
// Assumptions	5
assume <a1> (time_s = time_s.offset(by:</a1>	6
-1).defaults(to: 0)	
and time_us > time_us.offset(by: -1).defaults(to: 0))	7
and (time_s > time_s.offset(by: -1).defaults(to: 0)	8
or time_us > time_us.offset(by: -1).defaults(to:	9
0))	
// Exemplary State Statistics	10
<pre>output time := cast(time_s) + cast(time_us) / 1000000.0</pre>	11
output correct_vel := abs(vel) < 0.3	12
<pre>output cur_state := if correct_vel then</pre>	13
if cur_state.offset(by: -1).defaults(to: 0) = 0 then 1	14
else 2 else 0	
<pre>output start_interval := cur_state = 2</pre>	15
<pre>output interval_start := if start_interval then</pre>	16
interval_start.offset(by: -1).defaults(to: 0.0) else	
time	
trigger start_interval "Interval started!"	17
output end_interval := cur_state.offset(by:	18
-1).defaults(to: 0) > 0 and !correct_vel and	
time_since_start > 5.0	
trigger end_interval "Interval ended!"	19
<pre>output time_since_start := time -</pre>	20
interval_start.offset(by: -1).defaults(to: 0.0)	
// Assertions	21
assert <a1> !(start_interval and end_interval)</a1>	22
and time_since_start > 0.0	23

A.6 Specification: *ctrl_output*

import math	1
input time_s: UInt64	2
input time_us: UInt64	3
input vel_x: Float64	4
input vel_y: Float64	5
input vel_z: Float64	6
input fuel: Float64	7
input power: Float64	8
input vel_r_x: Float64	9
input vel_r_y: Float64	10

input vel_r_z: Float64	11
// Assumptions	12
assume $\langle a_1 \rangle$ trace_pos $\rangle = 0$	13
$(0.1) \leq 0.1$	14
and time – time.offset(by: -1).defaults(to: time – 0.1) > 0.0	15
assume <a2> power > 0.0</a2>	16
and power <= power.offset(by: -1).defaults(to:	17
power)	
and fuel > 0.0 and fuel < fuel.offset(by: -1).defaults(to: fuel + 0.1)	18
and $(\text{time}_s = \text{time}_s.\text{offset}(\text{by:} -1).\text{defaults}(\text{to:} 0)$	19
and time_us > time_us.offset(by: -1).defaults(to: 0)) and (time_s > time_s offset(by: -1) defaults(to: 0)	20
and $(\text{unne}_s > \text{unne}_s, \text{onset}(by; -1), \text{defaults}(to; 0)$	$\frac{21}{22}$
(0)) 0)	
// Frequency computations	23
<pre>output time := cast(time_s) + cast(time_us) / 1000000.0</pre>	24
<pre>output start_time := if time.offset(by: -1).defaults(to:</pre>	25
-1.0 == -1.0 then time else start_time.offset(by:	
-1).defaults(to: -1.0)	26
output flight_time := time - start_time	26
vel v or vel z or fuel or power or vel r x or	21
vel_y of vel_z of fuel of power of vel_1_x of vel r v or vel r z :=	
trace_pos.offset(by: -1).defaults(to: -1) + 1	
output frequency := 1.0 / (time – time.offset(by:	28
-1).defaults(to: time - 0.0001)) // major	
improvement	
output freq_sum := freq_sum.offset(by:	29
-1). defaults(to: 0.0) + frequency output freq avg := freq sum / cast(trace nos+1)	30
output freq_avg := if frequency > freq_max.offset(by:	31
-1).defaults(to: frequency) then frequency else	01
freq_max.offset(by: -1).defaults(to: frequency)	
<pre>output freq_min := if frequency < freq_min.offset(by:</pre>	32
-1).defaults(to: frequency) then frequency else	
freq_min.offset(by: -1).defaults(to: frequency)	
// Exemplary phase detection	33
vel_z*vel_z)	34
output velocity max := if reset max.offset(by:	35
-1).defaults(to: false) then velocity else max(
velocity, velocity_max.offset(by: -1).defaults(to:	
0.0))	
output velocity_min := if reset_max.offset(by:	36
-1).defaults(to: faise) then velocity else min(
0.0)	27
output un_inax := $abs(velocity_inax - velocity_min)$	31
output reset time := if reset max or trace $pos = 0$ then	30
time else reset time.offset(by: -1).defaults(to:	57
0.0)	
<pre>output unchanged := if reset_max.offset(by:</pre>	40
-1).defaults(to: false) then 0 else	
unchanged.onset(by: -1).defaults(to: 0) + 1 trianger unchanged = 150 "Phase detected!"	11
// Statistics	42
11 Statisties	

<pre>output vel_dev := abs(vel_r_x-vel_x) +</pre>	43
$abs(vel_r_y-vel_y) + abs(vel_r_z-vel_z)$	
<pre>output dev_sum := vel_dev + dev_sum.offset(by:</pre>	44
<pre>output vel_av := dev_sum / cast((trace_pos+1)*3)</pre>	45
<pre>output worst_dev_pos := if worst_dev.offset(by:</pre>	46
-1).defaults(to: vel_dev -1.0) < vel_dev then	
trace_pos else worst_dev_pos.offset(by:	
-1).defaults(to: 0)	
<pre>output worst_dev := if worst_dev.offset(by:</pre>	47
-1).defaults(to: vel_dev -1.0) < vel_dev then	
<pre>vel_dev else worst_dev.offset(by: -1).defaults(to:</pre>	
0.0)	
output start_fuel := start_fuel.offset(by:	48
-1).defaults(to: fuel)	
output fuel_level := 1 - (start_fuel - fuel) / start_fuel	49
output fuel_half := fuel_level < 0.50	50
output fuel_warning := fuel_level < 0.25	51
output fuel_danger := fuel_level < 0.10	52
output start_power := start_power.onset(by:	53
-1).defaults(to: power)	54
(start_power_p_consumed .= ((start_power - power))	54
trigger once fuel half "INFO: Fuel level is half	55
reduced"	55
trigger once fuel warning "WARNING: Fuel level is	56
below 25%"	00
trigger once fuel danger "DANGER: Fuel level is	57
below 10%"	
trigger_once power_p_consumed > 0.50 "Power below	58
half capacity"	
trigger_once power_p_consumed > 0.75 "Power below	59
quarter capacity"	
trigger_once power_p_consumed > 0.90 "Urgent:	60
Recharge Power!"	(1
// Assertions	61
assert <a1> time.onset(by: -1).defaults(to: -1.0) <</a1>	62
unie and start time start time affect(bu	62
and start_time == start_time.onset(by: -1) defaults(to: start_time)	05
and flight time \geq flight time offset(by:	64
-1) defaults(to: 0.0)	04
assert $< a^2 > reset$ time $> = 0.0$	65
and start fuel \geq fuel and start power \geq power	66
and (!fuel_half.offset(by: -1).defaults(to: false) or	67
fuel_half)	
and (!fuel_warning.offset(by: -1).defaults(to:	68
false)or fuel_warning)	
and (!fuel_danger.offset(by: -1).defaults(to: false) or	69
fuel_danger)	
and power_p_consumed >=	70
power_p_consumed.offset(by: -1).defaults(to:	
power_p_consumed)	

A.7 Specification: *mm_output_*1

import math	1
input stateID_SC: UInt64	2

// Assumptions	3
assume <a1> trace_pos >= 0</a1>	4
<pre>// Exemplary state transition analysis</pre>	5
output trace_pos @ stateID_SC := trace_pos.offset(by:	6
-1).defaults(to: -1) + 1	
output change_state := if trace_pos = 0 then false	7
<pre>else stateID_SC != stateID_SC.offset(by:</pre>	8
-1).defaults(to: 0)	
output transitions := if	9
<pre>stateID_SC.offset(by:-1).defaults(to: 0) = 0 then</pre>	
stateID_SC == 1	
else if stateID_SC.offset(by: -1).defaults(to: 0) == 1	10
then stateID_SC == 1 or stateID_SC == 2	
else if stateID_SC.offset(by: -1).defaults(to: 0) == 2	11
then stateID_SC == 1 or stateID_SC == 3	
else if stateID_SC.offset(by: -1).defaults(to: 0) == 3	12
then stateID_SC == 3	
else false	13
output invalid_transitions := change_state and	14
!transitions	
trigger invalid_transitions "Invalid state transition"	15
// Assertions	16
assert <a1> invalid_transitions or</a1>	17
!(stateID_SC.offset(by: -1).defaults(to: 0) != 0 and	18
stateID_SC = 0)	
assert <a2> (stateID_SC == 1 or stateID_SC == 2 or</a2>	19
stateID_SC == 3)	
or !(stateID_SC.offset(by: -2).defaults(to: 0) = 1	20
and transitions.offset(by: -1).defaults(to: false) and	21
transitions)	

A.8 Specification: *mm_output_*2

import math	1
input time_s: UInt64	2
input time_us: UInt64	3
input stateID_SC: Int64	4
input OnGround: UInt64	5
// Assumptions	6
assume <a1> trace_pos >= 0</a1>	7
and time - time.offset(by: -1).defaults(to: time -	8
$(0.1) \le 0.1$	
and time - time.offset(by: -1).defaults(to: time -	9
0.1) > 0.0	
<pre>assume <a2> (time_s = time_s.offset(by:</a2></pre>	10
-1).defaults(to: 0)	
and time_us > time_us.offset(by: -1).defaults(to: 0))	11
and (time_s > time_s.offset(by: -1).defaults(to: 0)	12
or time_us > time_us.offset(by: -1).defaults(to:	13
0))	
// Frequency computations	14
output time := cast(time_s) + cast(time_us) / 1000000.0	15
<pre>output start_time := if time.offset(by: -1).defaults(to:</pre>	16
-1.0 == -1.0 then time else start_time.offset(by:	
-1).defaults(to: -1.0)	
<pre>output flight_time := time - start_time</pre>	17
<pre>output trace_pos @ time_s or time_us or stateID_SC</pre>	18
or OnGround := trace_pos.offset(by:	
-1).defaults(to: -1) + 1	

output frequency := 1.0 / (time – time.offset(by:	19
-1).defaults(to: time -0.0001))	
<pre>output freq_sum := freq_sum.offset(by:</pre>	20
-1).defaults(to: 0.0) + frequency	
output freq avg := freq sum / cast(trace pos+1)	21
// Phase Statistics	22
output change state := if trace $pos = 0$ then false	23
else stateID_SC != stateID_SC.offset(by:	24
-1) defaults(to: 0)	
trigger change state	25
output entrance time := if change state then time	26
else entrance_time offset(by: -1) defaults(to:	27
time)	21
output hover, end := change state and	28
stateID SC offset(by: 1) defaults(to: 1) = 4	20
state $D_SC.onset(0y, -1).defaults(to, -1) = 4$	20
time	29
time – entrance_time.onset(by: -1).defaults(to:	30
	2.1
1) defende (co. 0.0) + h come com time.	31
-1).defaults(to: 0.0) + nover_cur_time	22
output nover_num_umes :=	32
$nover_num_times.onset(by: -1).defaults(to: 0) +$	
11 nover_end then 1 else 0	22
output nover_max_ume := max (33
$nover_max_time.onset(by: -1).defaults(to: 0.0),$	
nover_cur_time)	2.4
output nover_avg_time := if nover_num_times != 0	34
then nover_sum_time / cast(nover_num_times)	
	25
output landing_info := if change_state and stateID_SC	35
= 5 then 0.0 else time – entrance_time.offset(by:	
-1).defaults(to: time)	2.6
output landing_error := stateID_SC = 5 and OnGround	36
$!= 1$ and landing_into > 20.0	
// Assertions	37
assert <al> time.offset(by: -1).defaults(to: -1.0) <</al>	38
time	
and start_time == start_time.offset(by:	39
-1).defaults(to: start_time)	
and flight_time >= flight_time.offset(by:	40
-1).defaults(to: 0.0)	
<pre>assert <a2> time >= entrance_time and start_time <=</a2></pre>	41
entrance_time	
and hover_cur_time >= 0.0 and hover_max_time <=	42
flight_time	
assert <a3> !(landing_error and hover_end)</a3>	43
and $(!landing_error or landing_info > 0.0)$	44

A.9 Specification: contingency_output

input avgDist_laser: Float64	1
input actual_laser: Float64	2
input static_laser: Float64	3
<pre>input avgDist_optical:Float64</pre>	4
input actual_optical: Float64	5
input static_optical: Float64	6
// Assumptions	7
assume <a1> avgDist_laser >= 0.0 and actual_laser >=</a1>	8
0.0	

and static_laser >= 0.0 and avgDist_optical >= 0.0 and actual_optical >= 0.0 and static_optical >= 0.0 and (avgDist_laser + actual_laser + static_laser > 0.0) and (avgDist_optical + actual_optical + static_optical > 0.0)	9 10 11 12
// Trust computations	13
output rating_laser := 0.2 * static_laser + 0.4 *	14
actual_laser	
+ 0.4 * avgDist_laser	15
output rating_optical := 0.2 * static_optical + 0.4 *	16
actual_optical + 0.4 * avgDist_optical	
<pre>output trust_laser := rating_laser / (rating_laser +</pre>	17
rating_optical)	
output trust_optical := 1.0 - trust_laser	18
trigger trust_laser >= 0.5 "Trust in laser"	19
trigger trust_optical > 0.5 "Trust in optical sensor"	20
// Assertions	21
assert <a1> trust_laser ≠ trust_optical</a1>	22

A.10 Specification: health_output

import math	1
// average distance to the measured ostacle (range of	2
sight) using laser	
input avgDist_laser: Float64	3
// average distance to the measured ostacle (range of	4
sight) using camera	
<pre>input avgDist_optical: Float64</pre>	5
input vel: Float64	6
// Assumption	7
assume <a1> avgDist_laser <= 100.0 and</a1>	8
avgDist_laser >= 0.0 // both in m	
and avgDist_optical <= 50.0 and avgDist_optical >=	9
0.0 // both in m	
and $abs(vel) < 5.5 // in m/s$	1(
// Line of sight	11
<pre>output avgDst_dif := min(avgDist_laser,</pre>	12
avgDist_optical) – abs(vel)	
<pre>trigger avgDst_dif < 5.0 "WARNING: Dynamic</pre>	13
Velocity Limit reached"	
<pre>trigger avgDst_dif < 2.0 "ERROR: Abort mission."</pre>	14
// Assertions	15
assert <a1> avgDst_dif < 54.5 and avgDst_dif > -5.5</a1>	16

Funding Open Access funding enabled and organized by Projekt DEAL. This work was partially supported by the German Research Foundation (DFG) as part of the Collaborative Research Center Foundations of Perspicuous Software Systems (TRR 248, 389792660), by the European Research Council (ERC) Grant OSARES (No. 683300), and by the Aviation Research Program LuFo of the German Federal Ministry for Economic Affairs and Energy as part of "Volocopter Sicherheitstechnologie zur robusten eVTOL Flugzustandsabsicherung durch formales Monitoring" (No. 20Q1963C).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, 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 licence, and indicate if changes were made. The images or other third party material in this article are

included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence 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. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

- Baumeister, J., Finkbeiner, B., Schwenger, M., Torfah, H.: FFGA stream-monitoring of real-time properties. ACM Trans. Embed. Comput. Syst. 18(5s), 88 (2019). https://doi.org/10.1145/3358220
- Baumeister, J., Finkbeiner, B., Kruse, M., Schwenger, M.: Automatic optimizations for stream-based monitoring languages. In: Deshmukh, J., Ničković, D. (eds.) Runtime Verification, pp. 451–461. Springer, Cham (2020)
- Baumeister, J., Finkbeiner, B., Schirmer, S., Schwenger, M., Torens, C.: RTLola cleared for take-off: monitoring autonomous aircraft. In: Lahiri, S.K., Wang, C. (eds.) Computer Aided Verification, pp. 28–39. Springer, Cham (2020)
- Beckert, B., Hähnle, R., Schmitt, P.H. (eds.): Verification of Object-Oriented Software. The KeY Approach. LNCS, vol. 4334. Springer, Berlin (2007). https://doi.org/10.1007/978-3-540-69061-0
- Berry, G.: The foundations of ESTEREL. In: Proof, Language and Interaction: Essays in Honour of Robin Milner, pp. 425–454. MIT Press, Cambridge (2000)
- Bobot, F., Filliâtre, J.C., Marché, C., Paskevich, A.: Let's verify this with why3. Int. J. Softw. Tools Technol. Transf. 17, 709–727 (2015)
- Cluzeau, J.M., Henriquel, X., van Dijk, L., Gronskiy, A.: Concepts of design assurance for neural networks (CoDANN). Tech. Rep., EASA European Union Aviation Safety Agency (Mar 2020)
- D'Angelo, B., Sankaranarayanan, S., Sanchez, C., Robinson, W., Finkbeiner, B., Sipma, H.B., Mehrotra, S., Manna, Z.: Lola: runtime monitoring of synchronous systems. In: 12th International Symposium on Temporal Representation and Reasoning (TIME'05), pp. 166–174 (2005). https://doi.org/10.1109/TIME. 2005.26
- Dauer, J.C., Finkbeiner, B., Schirmer, S.: Monitoring with verified guarantees. In: Feng, L., Fisman, D. (eds.) Runtime Verification, pp. 62–80. Springer, Cham (2021)
- Faymonville, P., Finkbeiner, B., Schirmer, S., Torfah, H.: A streambased specification language for network monitoring. In: Falcone, Y., Sánchez, C. (eds.) Runtime Verification, pp. 152–168. Springer, Cham (2016)
- Finkbeiner, B., Oswald, S., Passing, N., Schwenger, M.: Verified rust monitors for lola specifications. In: Deshmukh, J., Ničković, D. (eds.) Runtime Verification, pp. 431–450. Springer, Cham (2020)
- Floyd, R.W.: Assigning meanings to programs. In: Program Verification, pp. 65–81. Springer, Dordrecht (1993). https://doi.org/10. 1007/978-94-011-1793-7_4
- Gautier, T., Le Guernic, P., Besnard, L.: SIGNAL: a declarative language for synchronous programming of real-time systems. In: Kahn, G. (ed.) Functional Programming Languages and Computer Architecture, pp. 257–277. Springer, Berlin (1987)
- Hagen, G., Tinelli, C.: Scaling up the formal verification of LUS-TRE programs with SMT-based techniques. In: 2008 Formal Methods in Computer-Aided Design, pp. 1–9 (2008). https://doi.org/10. 1109/FMCAD.2008.ECP.19
- Halbwachs, N., Caspi, P., Raymond, P., Pilaud, D.: The synchronous data flow programming language LUSTRE. Proc. IEEE 79(9), 1305–1320 (1991). https://doi.org/10.1109/5.97300

- Hoare, C.A.R.: An axiomatic basis for computer programming. Commun. ACM 12(10), 576–580 (1969). https://doi.org/10.1145/ 363235.363259
- Jagadeesan, L.J., Puchol, C., Von Olnhausen, J.E.: Safety property verification of ESTEREL programs and applications to telecommunications software. In: Wolper, P. (ed.) Computer Aided Verification, pp. 127–140. Springer, Berlin (1995)
- Leino, K.R.M.: Dafny: an automatic program verifier for functional correctness. In: Clarke, E.M., Voronkov, A. (eds.) Logic for Programming, Artificial Intelligence, and Reasoning, pp. 348–370. Springer, Berlin (2010)
- Müller, P., Schwerhoff, M., Summers, A.J.: Viper: a verification infrastructure for permission-based reasoning. In: Jobstmann, B., Leino, K.R.M. (eds.) Verification, Model Checking, and Abstract Interpretation, pp. 41–62. Springer, Berlin (2016)
- Nagarajan, P., Kannan, S.K., Torens, C., Vukas, M.E., Wilber, G.F.: ASTM F3269 – an Industry Standard on Run Time Assurance for Aircraft Systems. https://arc.aiaa.org/doi/abs/10.2514/6.2021-0525
- Nenzi, L., Bortolussi, L., Ciancia, V., Loreti, M., Massink, M.: Qualitative and quantitative monitoring of spatio-temporal properties. In: Bartocci, E., Majumdar, R. (eds.) Runtime Verification, pp. 21–37. Springer, Cham (2015)
- Pike, L., Goodloe, A., Morisset, R., Niller, S.: Copilot: a hard realtime runtime monitor. In: Barringer, H., Falcone, Y., Finkbeiner, B., Havelund, K., Lee, I., Pace, G., Roşu, G., Sokolsky, O., Tillmann, N. (eds.) Runtime Verification, pp. 345–359. Springer, Berlin (2010)

- Reactive Systems Group, C.: RTLola. https://github.com/ reactive-systems/RTLola-Frontend, https://github.com/reactivesystems/RTLola-Interpreter (2023)
- Reinbacher, T., Rozier, K.Y., Schumann, J.: Temporal-logic based runtime observer pairs for system health management of real-time systems. In: Ábrahám, E., Havelund, K. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, pp. 357–372. Springer, Berlin (2014)
- 25. Schirmer, S.: Runtime monitoring with Lola. Master's thesis, Saarland University (Dec. 2016)
- Schirmer, S., Torens, C., Adolf, F.: Formal monitoring of riskbased geofences. In: 2018 AIAA Information Systems-AIAA Infotech @ Aerospace (2018). https://arc.aiaa.org/doi/abs/10.2514/ 6.2018-1986
- Seto, D., Krogh, B., Sha, L., Chutinan, A.: The simplex architecture for safe online control system upgrades. In: Proceedings of the 1998 American Control Conference. ACC (IEEE Cat. No.98CH36207), vol. 6, pp. 3504–3508 (1998). https://doi.org/10.1109/ACC.1998. 703255
- Song, Y., Chin, W.N.: A synchronous effects logic for temporal verification of pure ESTEREL. In: Henglein, F., Shoham, S., Vizel, Y. (eds.) Verification, Model Checking, and Abstract Interpretation, pp. 417–440. Springer, Cham (2021)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.