



A Monitoring Tool for Linear-Time μ HML^{*}

Luca Aceto^{2,3} , Antonis Achilleos² , Duncan Paul Attard^{1,2}  ,
Léo Exibard² , Adrian Francalanza¹ , and Anna Ingólfssdóttir² 

¹ University of Malta, Msida, Malta {duncan.attard.01,afra1}@um.edu.mt

² Reykjavik University, Reykjavik, Iceland

{luca,antonios,duncanpa17,leoe,annai}@ru.is

³ Gran Sasso Science Institute, L'Aquila, Italy luca.aceto@gssi.it

Abstract. We present the implementation of a prototype tool that runtime checks specifications written in a maximally-expressive safety fragment of the linear-time modal μ -calculus called MAXHML. Our technical development is founded on previous results that give a compositional synthesis procedure for generating monitors from MAXHML formulae. This paper instantiates this synthesis to a first-order setting, where systems produce executions containing events that carry data. We augment the logic with predicates over data, and extend the synthesis procedure to generate executable monitors for Erlang, a general-purpose programming language. These monitors are instrumented via inlining to induce minimal runtime overhead. Our monitoring algorithm also maintains information, which it uses to explain how verdicts are reached.

Keywords: Runtime verification · Linear-time specifications · Monitor synthesis

1 Introduction

Runtime Verification (RV) [17,13,36,52] is a lightweight verification technique that dynamically checks the *current execution* to determine whether a System under Scrutiny (SuS) satisfies or violates some correctness stipulation. These stipulations are generally expressed using a specification logic to formally describe the behaviour the SuS should observe. RV synthesises specifications into *monitors*: computational entities that are instrumented with the SuS to analyse its execution (expressed as a trace of events) *incrementally*, and reach *verdicts* that cannot be retracted when observing future events. Fig. 1 depicts this set-up.

The vast majority of existing work and tooling efforts on RV focus on checking specifications that describe properties of *system executions* (Fig. 1a). Most of these studies are conducted in the context of temporal logics that are based

* Supported by the doctoral student grant (No: 207055) and the MoVeMnt project (No: 217987) of the Icelandic Research Fund, the BehAPI project funded by the EU H2020 RISE of the Marie Skłodowska-Curie action (No: 778233), the ENDEAVOUR Scholarship Scheme (Group B, national funds), and the MIUR project PRIN 2017FTXR7S IT MATTERS.

on LTL (e.g., [24,23,21,25,19,41,59,60,58]). Despite its widespread use, LTL has limited expressiveness. For instance, it cannot express properties such as ‘*every even position in the execution satisfies some proposition p* ’ [62,4].

The modal μ -calculus with a linear-time interpretation [49] has been shown to embed several other standard logics, including LTL, making it suitable to express a wider range of properties. Recent work [3,4] studies monitors for μ HML [7,51], a reformulation of the μ -calculus. One aspect that sets that work apart from the ones cited above is the modular approach the authors adopt in their technical development. Rather than redefining the semantics of the logic to assimilate the notion of monitoring verdicts, their study identifies *runtime monitorable* syntactic fragments of μ HML, delineating between its semantics on the one hand, and the operational semantics of monitors on the other. The authors define a synthesis procedure that generates *correct* monitors from these fragments. They also establish a correspondence between monitor acceptance (resp. rejection) verdicts and satisfactions (resp. violations) in the logic, and show that the fragments identified are *maximally-expressive, i.e.*, characterise all monitorable properties. This separation of concerns provides a principled approach to RV tool construction.

This paper presents the implementation of a prototype tool that builds on the theoretical foundations of [3,4]. It adopts the fragment MAXHML of μ HML that is used to specify safety properties on the *current system execution*. The study in [3,4] considers regular properties, which arguably limits its applicability to a broader setting where executions contain events that carry data [17,35]. We, therefore, lift the results of that study to a first-order setting, and extend the logic and synthesis procedure with predicates *over data*. Our adaptation of the monitor synthesis closely follows the one of [3,4], giving us high assurances that the corresponding monitors are correct. A facet that is often overlooked in RV is *verdict explainability*, where tools justify how monitoring judgements are reached [39]. Instantiations of this concept are commonplace in related fields. For example, model checking tools [48] produce diagnostic traces that explain why a model fails to satisfy some specification; in the same spirit, programming language frameworks capture runtime information and present it in the form of stack traces or core dumps. We take a first step in this direction, and engineer our monitoring algorithm to derive an explanation that is constructed using the monitor operational semantics of [3,4]. Since our prototype tool does not yet perform space optimisations for the purposes of explainable verdicts, this feature is intended for debugging or offline use.

In a parallel research direction, we study other monitorable μ HML fragments in a *branching-time* setting [37,38,2,1], where the logic describes properties about the *computation graph* of programs (Fig. 1b). Much of this body of work is concretised as `detectEr` [12,14,26,13], a runtime monitoring tool for concurrent Erlang programs. The material we propose in this paper complements the one in `detectEr`, contributing towards one tool that can runtime check specifications under their linear- and branching-time interpretations. Our contributions are:

- (i) We extend the logic MAXHML with data support, and show how this is used to specify properties on the current system execution, Sec. 2;

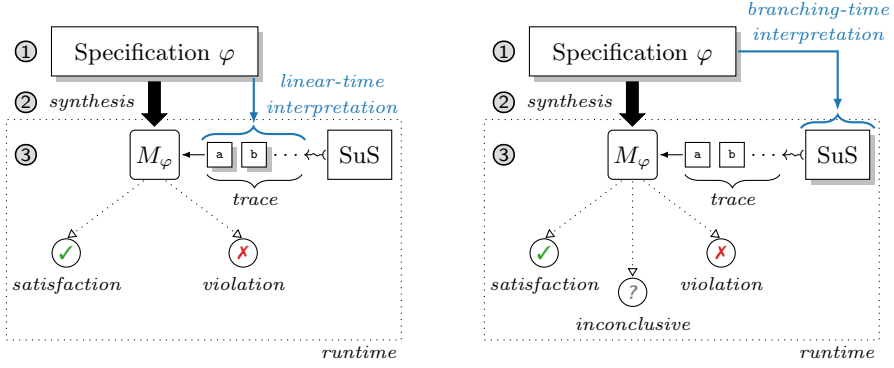


Fig. 1: RV of linear- and branching-time property φ by analysing the SuS trace

- (ii) adapt the monitor synthesis and operational semantics of [3,4] to enable monitors to reach verdicts based on the data carried by trace events, Sec. 3;
- (iii) discuss the challenges encountered when instantiating the synthesis in Sec. 3 to a general-purpose programming language, and overview the technique we use to instrument monitors that induce minimal runtime overhead, Sec. 4.

2 The Logic

We overview our chosen logic, MAXHML, a *maximally-expressive* syntactic fragment of μ HML used to describe safety properties of system executions [3]. It assumes a set of *external* actions $\alpha, \beta \in \text{ACT}$, together with a distinguished internal action $\tau \notin \text{ACT}$ that represents one *internal* step of computation. External actions range over values taken from some (potentially infinite) data domain, \mathbb{D} . Executions, also referred to as traces, are *infinite* sequences of external system actions that abstractly represent *complete* system runs. We reserve the metavariables $t, u \in \text{ACT}^\omega$ to represent infinite traces, and use αt to denote an infinite trace that starts with α and continues with t .

Fig. 2 shows our extension of MAXHML, called MAXHML^d, with predicates over data. Its syntax assumes a denumerable set of logical variables, $X, Y \in \text{LVAR}$. In addition to the standard Boolean constructs, the logic can express recursive properties as greatest fixed point formulae, $\max X.(\varphi)$, that bind the free occurrences of X in φ . The existential and universal modalities $\langle \mathbf{x}, e \rangle \varphi$ and $[\mathbf{x}, e] \varphi$ express the dual notions of *possibility* and *necessity* respectively. We augment these two modal constructs with *symbolic actions*, (\mathbf{x}, e) , to enable the reasoning on the data carried by external actions. Symbolic actions are pairs consisting of data variables, $x, y \in \text{DVAR}$, and *decidable* Boolean constraint expressions, $e, f \in \text{BEXP}$. Data variables range over the domain \mathbb{D} of data values, and bind

MAXHML^d Syntax

$$\varphi, \psi \in \text{MAXHML}^d ::= \text{tt} \mid \text{ff} \mid \langle \mathbf{x}, e \rangle \varphi \mid [\mathbf{x}, e] \varphi \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \max X.(\varphi) \mid X$$

MAXHML^d Semantics

$$\begin{aligned} \llbracket \text{tt}, \sigma \rrbracket &\triangleq \text{ACT}^\omega & \llbracket \text{ff}, \sigma \rrbracket &\triangleq \emptyset \\ \llbracket \langle \mathbf{x}, e \rangle \varphi, \sigma \rrbracket &\triangleq \{t \mid (\exists u. \exists \alpha. t = \alpha u \text{ and } e[\alpha/x] \Downarrow \text{true} \text{ and } u \in \llbracket \varphi[\alpha/x], \sigma \rrbracket)\} \\ \llbracket [\mathbf{x}, e] \varphi, \sigma \rrbracket &\triangleq \{t \mid (\forall u. \forall \alpha. (t = \alpha u \text{ and } e[\alpha/x] \Downarrow \text{true}) \text{ implies } u \in \llbracket \varphi[\alpha/x], \sigma \rrbracket)\} \\ \llbracket \varphi \vee \psi, \sigma \rrbracket &\triangleq \llbracket \varphi, \sigma \rrbracket \cup \llbracket \psi, \sigma \rrbracket & \llbracket \varphi \wedge \psi, \sigma \rrbracket &\triangleq \llbracket \varphi, \sigma \rrbracket \cap \llbracket \psi, \sigma \rrbracket \\ \llbracket \max X.(\varphi), \sigma \rrbracket &\triangleq \bigcup \{S \mid S \subseteq \llbracket \varphi, \sigma[X \mapsto S] \rrbracket\} & \llbracket X, \sigma \rrbracket &\triangleq \sigma(X) \end{aligned}$$

Fig. 2: Syntax and linear-time semantics for the logic MAXHML^d

the free occurrences of x in the expression e of the modality *and* in the continuation formula φ . The set BEXP, defined over \mathbb{D} and DVAR, consists of the usual Boolean operators \neg and \wedge , together with a set of relational operators that depends on \mathbb{D} , which we leave unspecified. For clarity, we omit writing the Boolean constraint expression e in modalities when $e = \text{true}$, and use *bold* lettering to identify binders in symbolic actions. In the sequel, the standard notions of open and closed expressions, and formula equality up to alpha-conversion are used. A formula is said to be *guarded* if every fixed point variable X appears within the scope of a modality that is itself in the scope of X . For example, $\max X.([\mathbf{x}] \text{ff} \wedge [\mathbf{y}] X)$ is guarded, as is $\max X.([\mathbf{x}]([\mathbf{y}] \text{ff} \wedge X))$, while $[\mathbf{x}] \max X.([\mathbf{y}] \text{ff} \wedge X)$ is not. Without loss of expressiveness [50], we assume all formulae to be guarded.

The linear-time interpretation of MAXHML^d is given by the denotational semantic function $\llbracket - \rrbracket$ that maps a formula to a *set* of traces. The function $\llbracket - \rrbracket$ uses valuations, $\sigma : \text{LVAR} \rightarrow 2^{\text{ACT}^\omega}$, to define the semantics inductively on the structure of formulae. The value $\sigma(X)$ is the set of traces that are assumed to satisfy X . In $\llbracket - \rrbracket$, modal formulae are interpreted w.r.t. symbolic actions. A symbolic action (\mathbf{x}, e) describes a *set* of external system actions. An action α is in this set when the data value it carries satisfies the Boolean constraint expression e that is instantiated with the *applied substitution* $[\alpha/x]$, i.e., $e[\alpha/x] \Downarrow \text{true}$ (see Fig. 2). The possibility formula $\langle \mathbf{x}, e \rangle \varphi$ denotes all the traces αu where α is in the action set (\mathbf{x}, e) and u satisfies the continuation $\varphi[\alpha/x]$. Dually, $[\mathbf{x}, e] \varphi$ denotes all the traces αu that, if prefixed by some α from the action set (\mathbf{x}, e) , u then satisfies $\varphi[\alpha/x]$. The set of traces satisfying the greatest fixed point formula $\max X.(\varphi)$ is the union of all the post-fixed point solutions, $S \subseteq \text{ACT}^\omega$, of the function induced by the formula φ . Since the interpretation of *closed* formulae does not depend on the environment σ , we may use $\llbracket \varphi \rrbracket$ in lieu of $\llbracket \varphi, \sigma \rrbracket$. A trace t satisfies (the closed) formula φ when $t \in \llbracket \varphi \rrbracket$, and violates φ when $t \notin \llbracket \varphi \rrbracket$.

To facilitate our exposition in this section and Sec. 3, we let $\mathbb{D} = \mathbb{Z}$, and fix the set of operators used in BEXP to \neg , \wedge and $=$. Sec. 4 considers the general case where the data carried by external actions can consist of *composite* data types. Henceforth, the terms *action* and *event* are used synonymously.




Fig. 3: Token server that issues integer identifier tokens to client programs

2.1 Trace Properties

Consider the model of a reactive token server, p_1 in Fig. 3, that issues client programs with identifier tokens that they use as an alias to write logs to a remote logging service. Clients request an identifier by sending the command 0, which the server then fulfils by replying with a new token, $n \in \mathbb{N}$. Since the server is itself a program that also uses the remote logging service, it is launched with its (reserved) identifier token 1. Fig. 3 shows that from its initial state p_1 , the token server either: (i) starts up with the token 1 and transitions to p_3 , where it waits for incoming client requests, or, (ii) fails to start and transitions with a status of -1 to the sink p_2 , thereafter exhibiting *undefined behaviour*. There are a number of properties we want *executions* of this token server to observe.

Example 1. One rudimentary property the current execution of server p_1 should uphold is that ‘no failure occurs at start up’. This safety requirement is expressed in terms of the MAXHML^d formula:

$$[\mathbf{x}, x = -1] \text{ff} \quad (\varphi_1)$$

The symbolic action $(\mathbf{x}, x = -1)$ defines the singleton set $\{-1\} \subset \mathbb{Z}$ of external system actions. Necessity modal formulae $[\mathbf{x}, e]\varphi$ state that, for *any* trace prefix α in the set defined by (\mathbf{x}, e) , the trace continuation u must then satisfy φ . However, *no* trace satisfies ff. This means that, in order for server traces *not* to violate formula φ_1 , they must start with actions $\alpha \notin \{-1\}$. The set of traces $1.(0.\mathbb{N})^\omega$ exhibited by p_1 satisfies this property, whereas $-1.\mathbb{Z}^\omega$ does not. ■

Example 2. Further to the stipulation of Ex. 1, we require that ‘the server is initialised with the identifier token 1’, expressed as:

$$[\mathbf{x}, x = -1] \text{ff} \wedge \langle \mathbf{x}, x = 1 \rangle \text{tt} \quad (\varphi_2)$$

The conjunct $[\mathbf{x}, x = -1] \text{ff}$ guards against traces of p_1 exhibiting failure when reading the token identifier; $\langle \mathbf{x}, x = 1 \rangle \text{tt}$ asserts that the trace exhibits 1 at start up, indicating a successful initialisation of the server. Formula φ_2 is satisfied exactly by server traces of the form $1.\mathbb{N}^\omega$. Note that the binders \mathbf{x} in $[\mathbf{x}, x = -1]$ and $\langle \mathbf{x}, x = 1 \rangle$ of φ_2 bind the variables x in *different* scopes. ■

The symbolic actions of Exs. 1 and 2 define sets of external actions w.r.t. literal values (e.g. $-1, 1$). More generally, action sets can be defined using constraint expressions that refer to other data variables within the same scope.

Example 3. Amongst the executions satisfying φ_2 are those where the server accidentally returns its identifier token 1 in reply to client requests. We therefore demand that ‘the server private token 1 is not leaked in client replies’. Formula φ_3 expresses this recursive property in a general way (note that Boolean constraint expressions $e = \text{true}$ are elided):

$$[\mathbf{x}] \max X.([\mathbf{y}]([\mathbf{z}, x = z] \text{ff} \wedge [\mathbf{z}, x \neq z] X)) \quad (\varphi_3)$$

The symbolic action $(\mathbf{x}, \text{true})$ in the first necessity defines the set of actions \mathbb{Z} . Its binder, \mathbf{x} , binds the variable x in $\max X.([\mathbf{y}]([\mathbf{z}, x = z] \text{ff} \wedge [\mathbf{z}, x \neq z] X))$. For some initial server action $\alpha \in \mathbb{Z}$, applying the substitution $[\alpha/x]$ to this continuation and unfolding the recursion variable once, gives the residual formula:

$$[\mathbf{y}]([\mathbf{z}, \alpha = z] \text{ff} \wedge [\mathbf{z}, \alpha \neq z] \max X.([\mathbf{y}]([\mathbf{z}, \alpha = z] \text{ff} \wedge [\mathbf{z}, \alpha \neq z] X)) \quad (\varphi_3')$$

The necessity $[\mathbf{y}]$ maps \mathbf{y} to the second server action β in the trace, *i.e.*, $[\beta/\mathbf{y}]$. Applying $[\beta/\mathbf{y}]$ to $[\mathbf{z}, \alpha = z] \text{ff}$ and $[\mathbf{z}, \alpha \neq z] \max X.([\mathbf{y}]([\mathbf{z}, \alpha = z] \text{ff} \wedge [\mathbf{z}, \alpha \neq z] X))$ leaves both sub-formulae unchanged, since \mathbf{y} binds no variables. For the third server action γ , the modalities $[\mathbf{z}, \alpha = z]$ and $[\mathbf{z}, \alpha \neq z]$ map \mathbf{z} to γ . Formula φ_3 is violated, ff , when the constraint $\alpha = z[\gamma/z]$ holds, *i.e.*, $\alpha \in \{n \in \mathbb{Z} \mid \alpha = \gamma\}$.

Crucially, a *fresh* scope for data variables is created upon each unfolding of X , such that \mathbf{y} and \mathbf{z} can be mapped to new values. By contrast, the value in \mathbf{x} is substituted for *once* in φ_3' and remains fixed when X is unfolded. Concretely, formula φ_3 compares actions at *every* odd position in the trace against the one at the head. Interpreting φ_3 over traces that the token sever exhibits on *successful* initialisation ensures that in particular, $1.(0.\{n \in \mathbb{N} \mid n \neq 1\})^*. (0.1).\mathbb{N}^\omega$ are violating. We remark that this property is *not* expressible in LTL. ■

3 Monitor Synthesis

The logic MAXHML^d is interpreted over infinite traces that represent *complete* system runs (refer to Sec. 2). In (online) RV where the SuS is *reactive*, obtaining complete runs is typically not possible [17,52], since the current trace corresponds to a prefix that is incrementally extended as the system execution unfolds. The notions of *good* and *bad prefixes* for monitorable properties provide sufficient evidence to determine acceptance or rejection. Informally, a good (resp. bad) prefix is a finite trace for which *every* extension satisfies (resp. violates) a property φ [10,24]. Monitors capture this principle through *irrevocable verdicts* that, once reached, cannot be retracted when observing future events.

Monitors may be viewed as processes via the syntax given in Fig. 4. This syntax differs from its regular counterpart of [3,4] in that it augments the prefixing construct with symbolic actions, (\mathbf{x}, e) . Besides the prefixing, external choice, and recursion constructs of CCS [54], the syntax of Fig. 4 includes disjunctive, \oplus , and conjunctive, \otimes , *parallel composition*. We use the symbol \odot to refer to both \oplus and \otimes when needed. Monitor verdict states, $v \in \text{VRD}$, are expressed as *yes* and *no*, respectively denoting *acceptance* and *rejection*.

Monitor Syntax

$$\begin{aligned}
m, n \in \text{MON} &::= v & | & (\mathbf{x}, e).m & | & m + n & | & m \oplus n & | & m \otimes n & | & \text{rec } X.m & | & X \\
v \in \text{VRD} &::= \text{yes} & | & \text{no}
\end{aligned}$$

Monitor Small-Step Semantics

$$\begin{aligned}
\text{MVRD} &\frac{}{v \xrightarrow{\alpha} v} & \text{MACT} &\frac{e[\alpha/x] \Downarrow \text{true}}{(\mathbf{x}, e).m \xrightarrow{\alpha} m[\alpha/x]} & \text{MCHSL} &\frac{m \xrightarrow{\alpha} m'}{m + n \xrightarrow{\alpha} m'} \\
\text{MTAUL} &\frac{m \xrightarrow{\tau} m'}{m \odot n \xrightarrow{\tau} m' \odot n} & \text{MPAR} &\frac{m \xrightarrow{\alpha} m' \quad n \xrightarrow{\alpha} n'}{m \odot n \xrightarrow{\alpha} m' \odot n'} & \text{MDISYL} &\frac{}{\text{yes} \oplus m \xrightarrow{\tau} \text{yes}} \\
\text{MDISNL} &\frac{}{\text{no} \oplus m \xrightarrow{\tau} m} & \text{MCONYL} &\frac{}{\text{yes} \otimes m \xrightarrow{\tau} m} & \text{MCONNL} &\frac{}{\text{no} \otimes m \xrightarrow{\tau} \text{no}} \\
\text{MREC} &\frac{}{\text{rec } X.m \xrightarrow{\tau} m[\text{rec } X.m/X]}
\end{aligned}$$

Monitor Synthesis

$$\begin{aligned}
\langle \text{tt} \rangle &= \text{yes} & \langle \text{ff} \rangle &= \text{no} \\
\langle (\mathbf{x}, e)\varphi \rangle &= (\mathbf{x}, e).\langle \varphi \rangle + (\mathbf{x}, \neg e).\text{no} & \langle [(\mathbf{x}, e)\varphi] \rangle &= (\mathbf{x}, e).\langle \varphi \rangle + (\mathbf{x}, \neg e).\text{yes} \\
\langle \varphi \vee \psi \rangle &= \langle \varphi \rangle \oplus \langle \psi \rangle & \langle \varphi \wedge \psi \rangle &= \langle \varphi \rangle \otimes \langle \psi \rangle \\
\langle \max X.\varphi \rangle &= \text{rec } X.\langle \varphi \rangle & \langle X \rangle &= X
\end{aligned}$$

Fig. 4: Syntax, synthesis, and small-step semantics for parallel monitors

Fig. 4 outlines the behaviour of monitors, where the transitions rules MREC , MCHSL , and its symmetric case MCHSR (omitted), are standard. Rule MACT describes the analysis monitors perform, where the binder \mathbf{x} in the symbolic action (\mathbf{x}, e) is mapped to an external system action α , yielding the substitution $[\alpha/x]$ that is applied to the Boolean constraint expression e . The monitor $(\mathbf{x}, e).m$ analyses α *only if* the instantiated constraint $e[\alpha/x]$ is satisfied, whereupon α is substituted for the *free* variable x in the body m . Verdict irrevocability is modelled by MVRD , where once in a verdict state v , any action can be analysed by monitors without altering v . Rule MPAR enables parallel sub-monitors to transition in lock-step when they analyse the *same* action α . The rest of the rules (omitting the obvious symmetric cases) cater for the internal reconfiguration of monitors. For instance, rules MDISYL and MDISNL state that in disjunctive parallelism, **yes** supersedes the verdicts of other monitors, whilst **no** does not affect the verdicts of other monitors; MCONYL and MCONNL express the dual case for parallel conjunctions. Finally, MTAUL and its symmetric analogue permit sub-monitors to execute internal reconfigurations independently.

Our adaptation $\langle - \rangle$ of the compositional synthesis procedure for regular monitors [3,4] is also given in Fig. 4. It generates monitors for $\varphi \in \text{MAXHML}^d$,

following the inductive structure of formulae. The translation for truth and falsehood, and the greatest fixed point and recursion variable constructs is direct; disjunction and conjunction are transformed to their parallel counterparts. Modal constructs are mapped to *deterministic* external choices, where the left summand handles the case where a system action α is in the set described by the symbolic action $\langle \mathbf{x}, e \rangle$, and the right summand, the case where α is *not* in this set. This embodies the duality of possibility and necessity: when α is not in the action set $\langle \mathbf{x}, e \rangle$, the formula $\langle \mathbf{x}, e \rangle \varphi$ is violated, whereas $[\mathbf{x}, e] \varphi$ is trivially satisfied.

Example 4. The monitor m_2 synthesised from formula φ_2 is:

$$\begin{aligned} (\varphi_2) &= ([\mathbf{x}, x = -1] \text{ff} \wedge \langle \mathbf{x}, x = 1 \rangle \text{tt}) = ([\mathbf{x}, x = -1] \text{ff}) \otimes (\langle \mathbf{x}, x = 1 \rangle \text{tt}) \\ &= ((\mathbf{x}, x = -1). \text{no} + (\mathbf{x}, x \neq -1). \text{yes}) \otimes ((\mathbf{x}, x = 1). \text{yes} + (\mathbf{x}, x \neq 1). \text{no}) \quad (m_2) \end{aligned}$$

When analysing the server traces $-1. \mathbb{Z}^\omega$, monitor m_2 reduces to $\text{no} \otimes \text{no}$ via the rule MPAR. Its premises are obtained by applying the MCHS_L and MACT to the left sub-monitor, and MCHS_R and MACT to the right sub-monitor, giving:

$$(\mathbf{x}, x = -1). \text{no} + (\mathbf{x}, x \neq -1). \text{yes} \xrightarrow{-1} \text{no} \quad \text{and} \quad (\mathbf{x}, x = 1). \text{yes} + (\mathbf{x}, x \neq 1). \text{no} \xrightarrow{-1} \text{no}$$

Monitor $\text{no} \otimes \text{no}$ afterwards transitions internally, $\text{no} \otimes \text{no} \xrightarrow{\tau} \text{no}$, via either rule MCONN_L or MCONN_R. Analogously, m_2 reaches the verdict **yes** when analysing the server traces $1. \mathbb{N}^\omega$. Recall that when in a verdict state, the monitor can *always* analyse future events by virtue of MVRD, flagging the *same* outcome. The behaviour of monitor m_2 corresponds to the property that φ_2 describes (refer to [3,4] for details). ■

Example 5. Consider the recursive monitor m_3 synthesised from formula φ_3 :

$$(\mathbf{x}). \text{rec } X. (\mathbf{y}). \left(((\mathbf{z}, x = z). \text{no} + (\mathbf{z}, x \neq z). \text{yes}) \otimes ((\mathbf{z}, x \neq z). X + (\mathbf{z}, x = z). \text{yes}) \right) \quad (m_3)$$

For the server traces $1.0.2.0.1.(0.\mathbb{N})^\omega$, m_3 instantiates \mathbf{x} to the value 1 at the head, and applies the substitution $[1/x]$ to the residual monitor, giving:

$$\text{rec } X. (\mathbf{y}). \left(((\mathbf{z}, 1 = z). \text{no} + (\mathbf{z}, 1 \neq z). \text{yes}) \otimes ((\mathbf{z}, 1 \neq z). X + (\mathbf{z}, 1 = z). \text{yes}) \right) \quad (m'_3)$$

Hereafter, m'_3 unfolds continually, ensuring that no event carries the value 1 observed at the head of the trace. At every even position, \mathbf{y} is instantiated with 0, whereas the binders \mathbf{z} in each of the sub-monitors composed in parallel compare the value carried by events occurring at odd trace positions against 1. Monitor m'_3 reaches the verdict **no** via these reductions:

$$\begin{aligned} m'_3 &\xrightarrow{\tau} (\mathbf{y}). \left(((\mathbf{z}, 1 = z). \text{no} + (\mathbf{z}, 1 \neq z). \text{yes}) \otimes ((\mathbf{z}, 1 \neq z). m'_3 + (\mathbf{z}, 1 = z). \text{yes}) \right) \quad (m''_3) \\ &\xrightarrow{0} ((\mathbf{z}, 1 = z). \text{no} + (\mathbf{z}, 1 \neq z). \text{yes}) \otimes ((\mathbf{z}, 1 \neq z). m'_3 + (\mathbf{z}, 1 = z). \text{yes}) \quad (m'''_3) \\ &\xrightarrow{2} \text{yes} \otimes m'_3 \xrightarrow{\tau} m'_3 \xrightarrow{\tau} m''_3 \xrightarrow{0} m'''_3 \xrightarrow{1} \text{no} \otimes \text{yes} \xrightarrow{\tau} \text{no} \xrightarrow{n} \text{no} \xrightarrow{n} \dots \end{aligned}$$

For the non-rejecting server traces $1.(0.\{n \in \mathbb{N} \mid n \neq 1\})^\omega$, monitor m'_3 visits the state $\text{yes} \otimes m'_3$ indefinitely, where m'_3 supersedes the uninfluential verdict **yes** following the rule MCONY_L. ■

4 Implementation

We implement our tool in Erlang [11,27], a general-purpose programming language that adopts the *actor model* of concurrency [44,8]. In this model, processes communicate exclusively by addressing *asynchronous messages* to one another via their uniquely-assigned Process ID (PID). Besides sending and receiving messages, processes can also fork other processes. This concurrency paradigm is tailored to reactive systems, making it ideal for our study.

4.1 Refining the Model

We refine the logic of Sec. 2 and monitor model of Sec. 3 to fit the Erlang use-case, where data can consist of composite types, such as tuples and lists. Accordingly, we generalise our definition of external actions as follows. Let $l \in \mathcal{L}$ be a finite set of *action labels*, and d_1, d_2, \dots be data values taken from a set of data domains $\mathcal{D} = \bigcup_{i \in \mathbb{N}} \mathbb{D}_i$ (e.g. integers, PIDs, tuples, etc.). An external action α is a tuple, (l, d_1, \dots, d_n) , where the first element is the label, and d_1, \dots, d_n is the *data payload* carried by α . We use the notation $l(d_1, \dots, d_n)$ to write α .

Patterns, $p \in \text{PAT}$, are counterparts to external system actions. These are defined as tuples, $(l, \mathbf{x}_1, \dots, \mathbf{x}_n)$, where x_1, x_2, \dots are data variables ranging over \mathcal{D} . Our revised definition of symbolic actions in the modal constructs $\langle p, e \rangle \varphi$ and $[p, e] \varphi$ uses these patterns instead of variables (cf. Sec. 2). The binders $\mathbf{x}_1, \dots, \mathbf{x}_n$ in p bind the free occurrences of x_1, \dots, x_n in the Boolean constraint e , and in the continuation φ . We define the function, $\text{match}(p, \alpha)$, to handle *pattern matching*. This function returns a substitution, $\pi : \text{DVAR} \rightarrow \mathcal{D}$, that maps the variables in p to the corresponding data values in the payload carried by α , when the shape of the pattern matches that of the action, or \perp if the match is unsuccessful. Analogous to the symbolic actions of Sec. 2, (p, e) describes a set of actions: an action α is in this set if (i) the pattern match succeeds, i.e., $\text{match}(p, \alpha) = \pi$, and (ii) the *instantiated* Boolean constraint expression $e\pi$ holds.

To instantiate our tool to Erlang, we use the action label set $\mathcal{L} = \{\rightarrow, \leftarrow, \star, !, ?\}$, that captures the lifecycle of, and interaction between processes. The *fork* action, \rightarrow , is exhibited by a process when it creates a child; its dual, \leftarrow , is exhibited by the child process upon *initialisation*. An *error* action, \star , signals abnormal process behaviour; *send* and *receive*, respectively $!$ and $?$, denote interaction. Tbl. 1 details the actions related to these labels, and the data payload they carry.




Fig. 5: Client-server interaction of the Erlang token server implementation

Action α	Action pattern p	Variables	Description
<i>fork</i>	$\rightarrow(x_1, x_2, y_1, y_2, y_3)$	x_1	PID of the parent process forking x_2
<i>initialise</i>	$\leftarrow(x_2, x_1, y_1, y_2, y_3)$	x_2	PID of the child process forked by x_1
		y_1, y_2, y_3	Function signature forked by x_1
<i>error</i>	$\star(x_1, y_1)$	x_1	PID of the erroneous process
		y_1	Error datum, <i>e.g.</i> error reason, <i>etc.</i>
<i>send</i>	$!(x_1, x_2, y_1)$	x_1	PID of the process sending the message
		x_2	PID of the recipient process
		y_1	Message datum, <i>e.g.</i> integer, tuple, <i>etc.</i>
<i>receive</i>	$?(x_2, y_1)$	x_2	PID of the recipient process
		y_1	Message datum, <i>e.g.</i> integer, tuple, <i>etc.</i>

Tbl. 1: Actions capturing the behaviour exhibited by Erlang processes

Our token server of Fig. 3 is readily translatable to Erlang, as shown in Fig. 5. The server starts when its main function, `lp`, in the Erlang module `ts` is invoked, state p_1 . From p_1 , it transitions to p_3 , exhibiting the initialisation event $\leftarrow(PID_s, PID_p, ts, lp, 1)$; the placeholders PID_s and PID_p respectively denote the PID values of the token server process and of the parent process forking the server. At p_3 , the server accepts client requests, consisting of the tuple $\{PID_c, 0\}$, where PID_c denotes the PID of the client, and 0 is the command requesting a new token. From state p_4 , the server replies with n , and transitions back to p_3 . This client-server interaction results in the server events $?(PID_s, \{PID_c, 0\})$ and $!(PID_s, PID_c, n)$. When the server fails at startup, it exhibits abnormal behaviour, shown as the error events $\star(PID_s, -1)$ and $\star(PID_s, m)$.

Example 6. Formula φ_2 w.r.t. the Erlang server of Fig. 5 is expressed as follows:

$$[\star(x_1, x_2), x_2 = -1] \text{ff} \wedge \langle \leftarrow(x_1, x_2, x_3, x_4, x_5), x_5 = 1 \rangle \text{tt} \quad (\varphi_4)$$

The patterns in the left and right conjuncts of φ_4 match the error and initialisation events. When p_1 exhibits an error at start up, $\text{match}(\star(x_1, x_2), \star(PID_s, -1))$, yields the substitution $\pi = [PID_s/x_1, -1/x_2]$, and the instantiated Boolean constraint $(x_2 = -1)\pi$ holds. For the same event, $\text{match}(\leftarrow(x_1, x_2, x_3, x_4, x_5), \star(PID_s, -1)) = \perp$ in the right conjunct, leading to a violation of formula φ_4 . The reverse argument applies for when p_1 loads successfully, where φ_4 is satisfied. In φ_4 , the pattern variables x_1 in $\star(x_1, x_2)$, and x_1, x_2, x_3, x_4 in $\leftarrow(x_1, x_2, x_3, x_4, x_5)$ are *redundant*.

$$[\leftarrow(_, _, _, _, x_5)] \max X \left([_] ([!(_, _, z_3), x_5 = z_3] \text{ff} \wedge [!(_, _, z_3), x_5 \neq z_3] X) \right) \quad (\varphi_5)$$

Formula φ_5 restates φ_3 with pattern matching. It uses the ‘don’t care’ pattern $_$, that matches *arbitrary* values, eliding redundant patterns and variables. ■

4.2 The Monitor Synthesis

Our synthesis from MAXHML^d specifications to executable Erlang monitors follows that of Fig. 4. Fig. 6 omits the cases for the falsity, necessity and conjunction constructs, as these are analogous to the ones for tt , $\langle p, e \rangle \varphi$ and $\varphi \vee \psi$. The translation from specifications to monitors is executed in three stages. First, a formula is parsed into its equivalent Abstract Syntax Tree (AST). This is then passed to the code generator that visits each of its nodes, mapping it to a *monitor description* as per the rules of Fig. 6. The monitor description is encoded as an Erlang AST to simplify its handling. In the final stage, this AST is processed by the Erlang compiler to emit the monitor source code or a BEAM [27] executable.

In this definition of $\langle - \rangle$, tt (resp. ff) is translated to the Erlang *atom yes* (resp. *no*) that indicates acceptance (resp. rejection). The remaining cases generate Erlang tuples whose first element, called the *tag*, is an atom that identifies the kind of monitor. Disjunctions (resp. conjunctions) are translated to the tuple tagged with *or* (resp. *and*), combining two sub-monitor descriptions. Greatest fixed point constructs, $\max X.(\varphi)$, are mapped to *rec* tuples consisting of *named* functions, $\text{fun } X() \rightarrow \langle \varphi \rangle \text{ end}$, that can be referenced by $\langle X \rangle$. Modal constructs are synthesised as a choice with *left* and *right* actions. An action tuple, *act*, combines a *predicate* function and an associated *monitor body* that is unfolded when the predicate is *true*. The predicate function encodes the pattern matching *and* Boolean constraint evaluation as one operation, using two *clauses*. Its first clause, $\text{fun}(p) \text{ when } e$, tests the constraint e w.r.t. the variables in the pattern p that become *dynamically* instantiated with the data values carried by an action α at runtime. The second catch-all clause $(_)$ covers the remaining cases, namely when: (i) either the action under analysis fails to match the pattern, or, (ii) the pattern matches *but* the Boolean constraint does *not* hold. For the left action, the predicate clause $\text{fun}(p) \text{ when } e$ returns *true* when the pattern match and guard test succeed, and *false* otherwise, *i.e.*, $(_)$. This condition is inverted for the right action, modelling cases (i) and (ii) just described. Our encoding of the aforementioned predicate in terms of Erlang function clauses spares us from implementing the pattern matching and constraint evaluation mechanism.

$$\begin{aligned}
 \langle \text{tt} \rangle &= \text{yes} & \langle \varphi \vee \psi \rangle &= \{\text{or}, \langle \varphi \rangle, \langle \psi \rangle\} \\
 \langle \max X.(\varphi) \rangle &= \{\text{rec}, \text{fun } X() \rightarrow \langle \varphi \rangle \text{ end}\} & \langle X \rangle &= \{\text{rec}, X\} \\
 \langle \langle p, e \rangle \varphi \rangle &= \left\{ \begin{array}{l} \{\text{chs}, \overbrace{\{\text{act}, \text{fun}(p) \text{ when } e \rightarrow \text{true}; (_) \rightarrow \text{false} \text{ end},}^{\text{predicate}} \\ \text{fun}(p) \rightarrow \langle \varphi \rangle \text{ end}\}, \\ \{\text{act}, \text{fun}(p) \text{ when } e \rightarrow \text{false}; (_) \rightarrow \text{true} \text{ end}, \\ \underbrace{\text{fun}(_) \rightarrow \text{no} \text{ end}}_{\text{monitor body}} \} \end{array} \right\} \begin{array}{l} \text{left action} \\ \text{right action} \end{array}
 \end{aligned}$$

Fig. 6: Translation from MAXHML^d formulae to Erlang code (excerpt)

```

1 def DERIVEACT( $\alpha, mon$ )
2 match  $mon$  do
3   case  $yes \vee no$ 
4     print 'Verdict reached'
5   case  $\{act, Pred, m\}$ 
6     return  $m(\alpha)$  # Apply m to event  $\alpha$ 
7   case  $\{chs, m, n\}$ 
8     if  $HOLDS(\alpha, m) \wedge \neg HOLDS(\alpha, n)$ 
9       return  $DERIVEACT(\alpha, m)$ 
10    else if  $\neg HOLDS(\alpha, m) \wedge HOLDS(\alpha, n)$ 
11      return  $DERIVEACT(\alpha, n)$ 
12    end if
13  case  $\{Op, m, n\} \wedge Op \in \{or, and\}$ 
14     $m' = DERIVEACT(\alpha, m)$ 
15     $n' = DERIVEACT(\alpha, n)$ 
16    return  $\{Op, m', n'\}$ 
17 end def

```

Expect: *Monitor must be in ready state*

```

18 def ANALYSEACT( $\alpha, m$ )
19    $m' = DERIVEACT(\alpha, m)$ 
20   return  $REDUCETAU(m')$ 
21 end def

```

```

22 def DERIVETAU( $mon$ )
23 match  $mon$  do
24   case  $\{or, yes, m\}$  return  $yes$ 
25   case  $\{or, no, m\}$  return  $m$ 
26   case  $\{and, yes, m\}$  return  $m$ 
27   case  $\{and, no, m\}$  return  $no$ 
28   case  $\{rec, m\}$  return  $m()$  # Unfold
29   case  $\{Op, m, n\} \wedge Op \in \{or, and\}$ 
30     if  $m' = DERIVETAU(m) \wedge m' \neq \perp$ 
31       return  $m'$ 
32     else
33       return  $DERIVETAU(n)$ 
34     end if
35   case Otherwise return  $\perp$ 
36 end def

```

```

37 def REDUCETAU( $m$ )
38   if  $m' = DERIVETAU(m) \wedge m' \neq \perp$ 
39     return  $REDUCETAU(m')$ 
40   else
41     return  $m$  # No more  $\tau$  reductions
42   end if
43 end def

```

Alg. 1: Algorithm that reduces monitors following the small-step rules of Fig. 4

It also enables monitors to support most of the Erlang data types and its full range of Boolean constraint expression syntax [11]. For similar reasons, $(\langle p, e \rangle \varphi)$ encodes the monitor body as $\mathbf{fun}(p) \rightarrow (\varphi) \mathbf{end}$ to delegate scoping to the Erlang language. This facilitates our synthesis and optimises the memory management of monitors by offloading this aspect onto the language runtime.

4.3 The Monitoring Algorithm

The synthesis procedure of Fig. 4 generates monitors that can runtime check formulae in parallel against the same position in the trace via disjunctive and conjunctive parallel composition. Our tool is however engineered to *emulate* parallel monitors, rather than forking processes and delegate their execution to the Erlang runtime. While the latter method tends to simplify the synthesis and runtime monitoring, we adopt the *former* approach for two reasons:

- (i) Previous empirical evidence suggests that parallelising via processes may induce high overhead when the RV set-up is considerably scaled [15]. A process-free design may render this overhead more manageable [6,5].
- (ii) Emulating parallel monitors requires us to tease apart the synthesised monitor description from its operational semantics. By separating these two aspects, our monitoring algorithm can track the operational rules it applies to reduce the monitor state, and use these to justify how verdicts are reached.

Our monitoring algorithm (Alg. 1) takes a monitor description m generated by $\langle - \rangle$, and performs successive reductions by applying m to events from the trace until a verdict is reached. Simultaneously, the algorithm maintains all the possible *active* states of the monitor as this is evolved from one state to the next. Alg. 1 encodes this reduction strategy using a series of **case** statements (lines 2–16 and 23–35), following the operational semantics of Fig. 4. Each **case** maps the first part of a rule conclusion to a *pattern*, enabling the monitoring algorithm to unambiguously **match** the rule to apply. The body of **cases** consists of a **return** statement that corresponds to the outcome dictated by the rule. Rules with premises (*e.g.* MCHS_L , MPAR , *etc.*) are reduced *recursively* by reapplying rules until an axiom is met, whereas axioms (*e.g.* MVRD , MDISN_L , *etc.*) reduce immediately. For example, the pattern $\{\text{chs}, m, n\}$ on line 7 specifies that MCHS_L and MCHS_R only apply to monitors of the form $m+n$. Selecting whether to reduce the left or right sub-monitor by analysing α is delegated to the function HOLDS . This instantiates the predicate encoded in **act** tuples with the data from α (see Fig. 6), returning the result of the predicate test. When the condition $\text{HOLDS}(\alpha, m) \wedge \neg \text{HOLDS}(\alpha, n)$ is **true**, $m+n$ is reduced to m , equivalent to the application of MCHS_L ; the argument for MCHS_R is symmetric.

The function ANALYSEACT of Alg. 1 conducts the runtime analysis. It ensures that once an action is analysed, the monitor is left in a state where it is *ready* to analyse the next action. We implement this logic by organising the application of the operational rules of Fig. 4 into two functions, DERIVEACT and DERIVETAU , according to the kind of action used to reduce the monitor. DERIVEACT on line 19 reduces the monitor *once* by applying it to the action under analysis, yielding m' . Subsequently, REDUCETAU reapplies the function DERIVETAU until all the internal transitions of the monitor are exhausted (lines 38–42). The cases on lines 24–27, corresponding to the axioms MDISY_L , MDISN_L , MCONY_L , MCONN_L , terminate redundant monitor states, and may be seen as a form of *garbage collection*. Due to space constraints, DERIVETAU omits the cases symmetric to those of lines 24–27.

Every single application of DERIVEACT and DERIVETAU results in a derivation that shows how a monitor evolves from one state to the next according to the operational rules of Fig. 4. The function ANALYSEACT keeps a complete history of these derivations internally. Derivations are represented as trees, rooted at the conclusion and terminating at the axiom nodes, that for each step: (i) maintain the monitor state consisting of the substitution π , and, (ii) the name of the rule used to derive the step. Maintaining the variable-value mapping in π across derivation steps demands that we track the changes in these variables for every active monitor state. This includes accounting for different binding scopes, variable shadowing, and the creation of fresh scopes when recursion variables, X , are unfolded. Our Erlang implementation of the functions listed in Alg. 1 in the tool incorporate this described logic. We are aware that storing the complete history ultimately impacts the performance of the runtime analysis. Our tool compromises by offering two operating modes, **normal** and **debug**, where in the latter, the full derivation history is stored in memory for the purpose of explainability.

4.4 Monitor Instrumentation

Our tool leverages the existing inlining [34] mechanism implemented in `detectEr` to instrument the SuS with monitors. While this approach assumes access to the source code of the SuS, it has been shown to induce lower overhead, by contrast to its outline counterpart [33,32,5]. The tool provides the meta keyword `with`, to identify the SuS components against which MAXHML^d specifications are runtime checked. Readers are referred to [12] for more details.

5 Case Study

We show the usability of our tool by applying it to an off-the-shelf Erlang web-server called `Cowboy` [45]. It delegates its socket management to `Ranch` (a socket acceptor pool for TCP protocols [46]), but forwards incoming HTTP client requests to *protocol handlers* that are forked dynamically by the webserver to service requests independently. We use our tool to runtime check MAXHML^d specifications describing fragments of the interaction protocol between the `Cowboy` and `Ranch` components. Our aim is to: (i) demonstrate the *expressiveness* of our logic by capturing properties of real-world software, (ii) validate the *applicability* of our monitoring and instrumentation techniques to third-party code, namely to applications built on top of the Erlang OTP middleware libraries, and, (iii) explore the *utility* of explainable verdicts for diagnosing software issues.

We redesign the token server of Fig. 5 as a REST web service deployed on `Cowboy`. The server generates identifier tokens using one of two formats, UUIDs, or short alphanumeric strings. Clients request tokens by issuing a GET request with parameter, `type=uuid` or `type=short`, specifying the token format required. The web service offers a standard interface: (i) it returns HTTP 400 when the `type` parameter is omitted from the request, and, (ii) HTTP 500 when an unsupported `type` is used. We also simulate intermittent faults in the `Cowboy` components by *injecting* process crashes based on a fair Bernoulli trial [55].

For our case study, we consider a selection of properties describing the `Cowboy-Ranch` interaction protocol. One such property, φ_{rp} , concerns *Cowboy request processes* that service client requests. It states that in its (current) execution, ‘a request process does not issue HTTP responses with code 500, nor does it crash’.

$$\max X. \left(\begin{array}{l} [!(rprc, _, \{tag, code, \dots\}), tag = \mathbf{resp} \wedge code = 200] X \wedge \\ [!(rprc, _, \{tag, code, \dots\}), tag = \mathbf{resp} \wedge code = 500] \text{ff} \wedge \\ [\star(rprc, stat), stat = \mathbf{crash}] \text{ff} \end{array} \right) \quad (\varphi_{rp})$$

In φ_{rp} , the binders *tag* and *code* become instantiated with the atom `resp` designating a response message, and the HTTP code of the response returned to requesting clients. Besides ensuring that response messages sent by request processes do not contain the code 500, *i.e.*, $tag = \mathbf{resp} \wedge code = 500$, formula φ_{rp} also asserts that these processes do not *crash*, *i.e.*, $stat = \mathbf{crash}$. The binder *rprc*, referring to the request process PID, is included in φ_{rp} for clarity. While the monitor synthesised from φ_{rp} flags the corresponding rejection, the verdict

(alone) does not indicate the source of the error. This may suffice for verifying small-scale systems where errors are manually-trackable, but becomes impractical in realistic settings such as this case study. Our algorithm addresses this shortcoming by giving a justification showing how a monitor reaches its verdict.

6 Conclusion

This paper presents the implementation of a RV tool that runtime checks specifications written in a safety-fragment of the linear-time modal μ -calculus, called MAXHML, augmented with predicates over data. Our work builds on previous theoretical results for the regular setting [3,4] that give a compositional synthesis procedure which generates monitors from MAXHML formulae. We extend the logic, synthesis, and monitor operational semantics of [3,4] to enable the tool to handle events that carry data. We discuss the implementability of this synthesis procedure, and overview the approach our monitoring algorithm takes towards providing justifiable monitoring verdicts. Our tool is validated via a realistic case study that uses an off-the-shelf, third-party Erlang webserver. We show how our augmented logic flexibly expresses properties involving data, and argue for the utility of explainable verdicts for diagnosing software issues.

Related Work The synthesis procedure in this paper contrasts with another by a line of work that investigates the monitorable safety fragment of the branching-time modal μ -calculus, called sHML [37,38]. The latter synthesis generates monitors with non-deterministic behaviour that, while sufficient for the theoretical results required in *op. cit.*, may lead to missed detections in practice. An early materialisation of [37,38] as the tool `detectEr` [14,13,26,40] addresses this shortcoming by parallelising monitors using processes, enabling them to reach verdicts along all possible paths. While effective, this approach scales poorly [15]. Ongoing work on `detectEr` [12] indicates that a process-free approach could lead to more efficient runtime monitoring that scales considerably better [6,5].

There are other ways to monitoring systems with events that carry data besides the ones cited in Sec. 1 (see *e.g.*, [18,20,43,41,42,22,61]). One work that shares characteristics with ours is Parametric Trace Slicing (PTS) [29,57], where the global trace is projected into local sub-traces called *slices*, based on parametric specifications. These are properties specified in terms of *symbolic events* whose parameters are instantiated to values from events in the global trace. We use similar means to identify the SuS components to be instrumented, thus filtering out events and obtain trace slices (see Sec. 4.4). PTS is adopted by a number of RV tools that handle data (see *e.g.*, [9,53,28,47,31]), notably `MarQ` [56,16] for Java, and `Elarva` [30] for Erlang. `Elarva` takes a naïve strategy to PTS, relying on a central process to collect trace events that are demultiplexed between monitors to obtain slices. This makes it susceptible to single point of failures, and does not scale in practice. The design we use is more robust, as it directly instruments components of the SuS, giving us a modicum of fault containment when monitoring independently-executing components that can fail in isolation.

References

1. Aceto, L., Achilleos, A., Francalanza, A., Ingólfssdóttir, A.: A Framework for Parameterized Monitorability. In: FoSSaCS. LNCS, vol. 10803, pp. 203–220 (2018)
2. Aceto, L., Achilleos, A., Francalanza, A., Ingólfssdóttir, A., Kjartansson, S.Ö.: Determinizing Monitors for HML with Recursion. *JLAMP* **111**, 100515 (2020)
3. Aceto, L., Achilleos, A., Francalanza, A., Ingólfssdóttir, A., Lehtinen, K.: Adventures in Monitorability: From Branching to Linear Time and Back Again. *Proc. ACM Program. Lang.* **3**(POPL), 52:1–52:29 (2019)
4. Aceto, L., Achilleos, A., Francalanza, A., Ingólfssdóttir, A., Lehtinen, K.: An Operational Guide to Monitorability with Applications to Regular Properties. *Softw. Syst. Model.* **20**(2), 335–361 (2021)
5. Aceto, L., Attard, D.P., Francalanza, A., Ingólfssdóttir, A.: A Choreographed Outline Instrumentation Algorithm for Asynchronous Components. *CoRR* **abs/2104.09433** (2021)
6. Aceto, L., Attard, D.P., Francalanza, A., Ingólfssdóttir, A.: On Benchmarking for Concurrent Runtime Verification. In: FASE. LNCS, vol. 12649, pp. 3–23 (2021)
7. Aceto, L., Ingólfssdóttir, A., Larsen, K.G., Srba, J.: *Reactive Systems: Modelling, Specification and Verification*. Cambridge University Press (2007)
8. Agha, G., Mason, I.A., Smith, S.F., Talcott, C.L.: *A Foundation for Actor Computation*. *JFP* **7**(1), 1–72 (1997)
9. Allan, C., Avgustinov, P., Christensen, A.S., Hendren, L.J., Kuzins, S., Lhoták, O., de Moor, O., Sereni, D., Sittampalam, G., Tibble, J.: Adding Trace Matching with Free Variables to AspectJ. In: OOPSLA. pp. 345–364. ACM (2005)
10. Alpern, B., Schneider, F.B.: Defining Liveness. *Inf. Process. Lett.* **21**(4), 181–185 (1985)
11. Armstrong, J.: *Programming Erlang: Software for a Concurrent World*. Pragmatic Bookshelf (2007)
12. Attard, D.P., Aceto, L., Achilleos, A., Francalanza, A., Ingólfssdóttir, A., Lehtinen, K.: Better Late than Never or: Verifying Asynchronous Components at Runtime. In: FORTE. LNCS, vol. 12719, pp. 207–225. Springer (2021)
13. Attard, D.P., Cassar, I., Francalanza, A., Aceto, L., Ingólfssdóttir, A.: Introduction to Runtime Verification. In: *Behavioural Types: from Theory to Tools*, pp. 49–76. Automation, Control and Robotics, River (2017)
14. Attard, D.P., Francalanza, A.: A Monitoring Tool for a Branching-Time Logic. In: *RV*. LNCS, vol. 10012, pp. 473–481 (2016)
15. Attard, D.P., Francalanza, A.: Trace Partitioning and Local Monitoring for Asynchronous Components. In: SEFM. LNCS, vol. 10469, pp. 219–235 (2017)
16. Barringer, H., Falcone, Y., Havelund, K., Reger, G., Rydeheard, D.E.: Quantified Event Automata: Towards Expressive and Efficient Runtime Monitors. In: *FM*. LNCS, vol. 7436, pp. 68–84. Springer (2012)
17. Bartocci, E., Falcone, Y., Francalanza, A., Reger, G.: Introduction to Runtime Verification. In: *Lectures on RV*, LNCS, vol. 10457, pp. 1–33. Springer (2018)
18. Basin, D.A., Klaedtke, F., Müller, S., Zalinescu, E.: Monitoring Metric First-Order Temporal Properties. *J. ACM* **62**(2), 15:1–15:45 (2015)
19. Basin, D.A., Klaedtke, F., Zalinescu, E.: Failure-Aware Runtime Verification of Distributed Systems. In: *FSTTCS. LIPIcs*, vol. 45, pp. 590–603. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2015)
20. Basin, D.A., Klaedtke, F., Zalinescu, E.: Runtime Verification of Temporal Properties over Out-of-Order Data Streams. In: *CAV (1)*. LNCS, vol. 10426, pp. 356–376. Springer (2017)

21. Bauer, A., Falcone, Y.: Decentralised LTL Monitoring. *FMSD* **48**(1-2), 46–93 (2016)
22. Bauer, A., Küster, J., Vegliach, G.: The Ins and Outs of First-Order Runtime Verification. *Formal Methods Syst. Des.* **46**(3), 286–316 (2015)
23. Bauer, A., Leucker, M., Schallhart, C.: Comparing LTL Semantics for Runtime Verification. *J. Log. Comput.* **20**(3), 651–674 (2010)
24. Bauer, A., Leucker, M., Schallhart, C.: Runtime Verification for LTL and TLTL. *ACM Trans. Softw. Eng. Methodol.* **20**(4), 14:1–14:64 (2011)
25. Bonakdarpour, B., Fraigniaud, P., Rajsbbaum, S., Rosenblueth, D.A., Travers, C.: Decentralized Asynchronous Crash-Resilient Runtime Verification. In: *CONCUR. LIPIcs*, vol. 59, pp. 16:1–16:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016)
26. Cassar, I., Francalanza, A., Attard, D.P., Aceto, L., Ingólfssdóttir, A.: A Suite of Monitoring Tools for Erlang. In: *RV-CuBES. Kalpa Publications in Computing*, vol. 3, pp. 41–47 (2017)
27. Cesarini, F., Thompson, S.: *Erlang Programming: A Concurrent Approach to Software Development*. O’Reilly Media (2009)
28. Chen, F., Rosu, G.: Mop: An Efficient and Generic Runtime Verification Framework. In: *OOPSLA*. pp. 569–588 (2007)
29. Chen, F., Rosu, G.: Parametric Trace Slicing and Monitoring. In: *TACAS. LNCS*, vol. 5505, pp. 246–261 (2009)
30. Colombo, C., Francalanza, A., Gatt, R.: Elarva: A Monitoring Tool for Erlang. In: *RV. LNCS*, vol. 7186, pp. 370–374. Springer (2011)
31. Decker, N., Harder, J., Scheffel, T., Schmitz, M., Thoma, D.: Runtime Monitoring with Union-Find Structures. In: *TACAS. LNCS*, vol. 9636, pp. 868–884. Springer (2016)
32. Erlingsson, Ú.: *The Inlined Reference Monitor Approach to Security Policy Enforcement*. Ph.D. thesis, Cornell University (2004)
33. Erlingsson, Ú., Schneider, F.B.: SASI Enforcement of Security Policies: A Retrospective. In: *NSPW*. pp. 87–95 (1999)
34. Falcone, Y., Krstic, S., Reger, G., Traytel, D.: A Taxonomy for Classifying Runtime Verification Tools. In: *RV. LNCS*, vol. 11237, pp. 241–262 (2018)
35. Francalanza, A.: A Theory of Monitors. *Inf. Comput.* **281**, 104704 (2021)
36. Francalanza, A., Aceto, L., Achilleos, A., Attard, D.P., Cassar, I., Della Monica, D., Ingólfssdóttir, A.: A Foundation for Runtime Monitoring. In: *RV. LNCS*, vol. 10548, pp. 8–29 (2017)
37. Francalanza, A., Aceto, L., Ingólfssdóttir, A.: On Verifying Hennessy-Milner Logic with Recursion at Runtime. In: *RV. LNCS*, vol. 9333, pp. 71–86 (2015)
38. Francalanza, A., Aceto, L., Ingólfssdóttir, A.: Monitorability for the Hennessy-Milner Logic with Recursion. *FMSD* **51**(1), 87–116 (2017)
39. Francalanza, A., Cini, C.: Computer Says No: Verdict Explainability for Runtime Monitors using a Local Proof System. *J. Log. Algebraic Methods Program.* **119**, 100636 (2021)
40. Francalanza, A., Seychell, A.: Synthesising Correct Concurrent Runtime Monitors. *FMSD* **46**(3), 226–261 (2015)
41. Havelund, K., Peled, D.: Runtime Verification: From Propositional to First-Order Temporal Logic. In: *RV. LNCS*, vol. 11237, pp. 90–112. Springer (2018)
42. Havelund, K., Peled, D.: BDDs for Representing Data in Runtime Verification. In: *RV. LNCS*, vol. 12399, pp. 107–128. Springer (2020)

43. Havelund, K., Reger, G., Thoma, D., Zalinescu, E.: Monitoring Events that Carry Data. In: Lectures on Runtime Verification, LNCS, vol. 10457, pp. 61–102. Springer (2018)
44. Hewitt, C., Bishop, P.B., Steiger, R.: A Universal Modular ACTOR Formalism for Artificial Intelligence. In: IJCAI. pp. 235–245. William Kaufmann (1973)
45. Hoguin, L.: Cowboy (2020), <https://ninenines.eu>
46. Hoguin, L.: Ranch (2020), <https://ninenines.eu>
47. Jin, D., Meredith, P.O., Lee, C., Rosu, G.: JavaMOP: Efficient Parametric Runtime Monitoring Framework. In: ICSE. pp. 1427–1430 (2012)
48. Jr., E.M.C., Grumberg, O., Peled, D.A.: Model Checking. MIT Press (1999)
49. Kozen, D.: Results on the Propositional μ -Calculus. In: ICALP. LNCS, vol. 140, pp. 348–359 (1982)
50. Kupferman, O., Vardi, M.Y., Wolper, P.: An Automata-Theoretic Approach to Branching-Time Model Checking. J. ACM **47**(2), 312–360 (2000)
51. Larsen, K.G.: Proof Systems for Satisfiability in Hennessy-Milner Logic with Recursion. TCS **72**(2&3), 265–288 (1990)
52. Leucker, M., Schallhart, C.: A Brief Account of Runtime Verification. JLAP **78**(5), 293–303 (2009)
53. Meredith, P.O., Jin, D., Griffith, D., Chen, F., Rosu, G.: An Overview of the MOP Runtime Verification Framework. STTT **14**(3), 249–289 (2012)
54. Milner, R.: Communication and Concurrency. Prentice Hall (1989)
55. Papoulis, A.: Probability, Random Variables, and Stochastic Processes. McGraw Hill (1991)
56. Reger, G., Cruz, H.C., Rydeheard, D.E.: MarQ: Monitoring at Runtime with QEA. In: TACAS. LNCS, vol. 9035, pp. 596–610 (2015)
57. Reger, G., Rydeheard, D.E.: From First-Order Temporal Logic to Parametric Trace Slicing. In: RV. LNCS, vol. 9333, pp. 216–232. Springer (2015)
58. Scheffel, T., Schmitz, M.: Three-Valued Asynchronous Distributed Runtime Verification. In: MEMOCODE. pp. 52–61 (2014)
59. Sen, K., Vardhan, A., Agha, G., Rosu, G.: Efficient Decentralized Monitoring of Safety in Distributed Systems. In: ICSE. pp. 418–427 (2004)
60. Sen, K., Vardhan, A., Agha, G., Rosu, G.: Decentralized Runtime Analysis of Multithreaded Applications. In: IPDPS. IEEE (2006)
61. Stolz, V.: Temporal Assertions with Parametrized Propositions. J. Log. Comput. **20**(3), 743–757 (2010)
62. Wolper, P.: Temporal Logic Can be More Expressive. Inf. Control. **56**(1/2), 72–99 (1983)