

Synthesising correct concurrent runtime monitors

Adrian Francalanza · Aldrin Seychell

Published online: 20 November 2014
© Springer Science+Business Media New York 2014

Abstract This paper studies the correctness of automated synthesis for concurrent monitors. We adapt a subset of the Hennessy–Milner logic with recursion (a reformulation of the modal μ -calculus) to specify safety properties for Erlang programs. We also define an automated translation from formulas in this sub-logic to concurrent Erlang monitors that detect formula violations at runtime. Subsequently, we formalise a novel definition for monitor correctness that incorporates monitor behaviour when instrumented with the program being monitored. Finally, we devise a sound technique that allows us to prove monitor correctness in stages; this technique is used to prove the correctness of our automated monitor synthesis.

Keywords Runtime verification · Automated monitor synthesis · Monitor correctness · μ -calculus · Concurrency · Actors

1 Introduction

Runtime verification (RV), is a lightweight verification technique that may be used to determine whether the *current system run* respects a correctness property. Two requirements are crucial for the adoption of this technique. First, runtime *monitor overheads* need to be kept to a minimum so as not to degrade system performance. Second, instrumented monitors need to form part of the trusted computing base of the verification setup by adhering to an agreed notion of *monitor correctness*; amongst other things, this normally includes a guarantee that runtime checking corresponds (in some sense) to the property being checked for. Monitor overheads and correctness are occasionally conflicting concerns. For instance, in order to lower monitoring overheads, engineers increasingly use *concurrent monitors* [16,29,35] so as to benefit from the underlying parallel and distributed architectures. However concurrent

A. Francalanza (✉) · A. Seychell
Department of Computer Science, ICT, University of Malta, Msida, Malta
e-mail: adrian.francalanza@um.edu.mt

A. Seychell
e-mail: aldrin.seychell.08@um.edu.mt

monitors are also more susceptible to elusive errors such as non-deterministic behaviour, deadlocks or livelocks which may, in turn, affect their correctness.

Ensuring monitor correctness is, in general, non-trivial. One prominent obstacle is the fact that system properties are typically specified using one formalism, e.g. a high-level logic, whereas the respective monitors that check these properties are described using another formalism, e.g. a programming language. This, in turn, makes it hard to ascertain the semantic correspondence between the two descriptions. *Automated monitor synthesis* can mitigate this problem by standardising the translation from the property logic to the monitor formalism. It also gives more scope for a formal treatment of monitor correctness.

In this work, we investigate the correctness of *synthesised* monitors in a *concurrent setting*, whereby (i) the system being verified executes concurrently with the synthesised monitor (ii) the system and the monitor themselves consist of concurrent sub-components. The correctness of monitor synthesis has been studied previously by the seminal work of Geilen, [23], and (more formally) by subsequent work such as that of Sen et al. [34], and Bauer et al. [5]. Our approach differs from these studies in a number of respects. First, the aforementioned work abstracts away from the internal working of a system, representing it as a string of events/states (execution trace); this complicates reasoning about how monitors are instrumented wrt. an executing system. It also focusses on a logic that is readily amenable to runtime analysis, namely linear temporal logic (LTL) [13]. Moreover, it expresses synthesis in terms of *abstract* or *single-threaded* monitors—using either pseudocode [23] or automata [5,34]—executing wrt. such trace. By contrast, we strive towards a more intensional formal definition of online correctness for synthesised concurrent monitors whereby, for arbitrary property φ , the synthesised monitor M_φ running concurrently wrt. some system S (denoted as $S \parallel M_\varphi$) respects the following condition:

$$\text{Whenever } S \parallel M_\varphi \text{ executes to some } S' \parallel M' \quad \text{then} \quad \left(\begin{array}{l} \text{If the execution from } S \text{ to } S' \text{ violates } \varphi \text{ then} \\ \quad M' \text{ should consistently flag the violation} \\ \quad \text{and} \\ \text{If } M' \text{ flags a violation then the execution from } S \text{ to } S' \text{ should violate } \varphi \end{array} \right) \quad (1)$$

The setting of (1) brings to the fore additional issues concerning monitor correctness:

- (i) A property logic semantics may be defined over other computational entities apart from traces, i.e. the current *execution* in (1), such as the entire computational tree of a program. As a result, the logic semantics may not readily lend itself to the formulation of monitor correctness outlined in condition (1) above, which *only* requires monitor detection *whenever a violation occurs*. In general, a monitored system that violates a property according to the original logic semantics, may do so along one computational path but *not* along another; this is often the case for concurrent systems with multiple execution paths as a result of different thread interleavings scheduled at runtime.
- (ii) Concurrent monitors may also have multiple execution paths. Condition (1) thus requires stronger guarantees than those for single-threaded monitors so as to ensure that *all* these paths correspond to an appropriate runtime check of system property being monitored. Stated otherwise, correct concurrent monitors must *always/consistently* detect violations and flag them, irrespective of their runtime interleaving.
- (iii) Apart from the formal semantics of the source logic (specifying property φ), we also require a formal semantics for the target languages of *both* the system and the monitor executing in *side-by-side*, i.e. $S \parallel M_\varphi$. In most cases, the latter may not be available.

- (iv) Online monitor correctness needs to ensure that monitor execution cannot be interfered with by the system, and viceversa. Whereas adequate monitor instrumentation typically prevents direct interferences, condition (1) must consider *indirect* interferences such as system divergences [25,32], i.e. infinite internal looping making the system unresponsive, which may prevent the monitors from progressing.
- (v) Ensuring correctness along the lines of (1) can be quite onerous because *every* execution path of the monitor running concurrently with the monitored system, $S \parallel M_\phi$, needs to be analysed to ensure consistent detections along every thread interleaving. Consequently, one needs to devise scalable techniques facilitating monitor correctness analysis.

We conduct our study in terms of actor-based [26] concurrent monitors written in Erlang [3, 11], an industry-strength language for constructing fault-tolerant systems. To alleviate our technical development, we also restrict monitoring to systems written in the same language. We limit ourselves to a logic that describes *reactive properties*, i.e. *system interactions with the environment*, and focus on the synthesis of *asynchronous* monitors, performing runtime analysis through the Erlang virtual machine (EVM)'s tracing mechanism. Despite the typical drawbacks associated with asynchrony, e.g. late detections, our monitoring setup is in line with the asynchrony advocated by the actor concurrency model, which facilitates scalable coding techniques such as fail-fast design patterns [11]. Asynchronous monitoring has also been used in other RV tools, e.g. [14,17], and has proved to be less intrusive and easier to instrument than synchronous monitoring setups. It is also the more feasible alternative when monitoring distributed systems [16,20]. More importantly, though, it still allows us to investigate the main issues arising from the correctness setup outlined in (1), and we expect most of the issues investigated to carry over in a straightforward fashion to purely synchronous settings.

As an expository logic for describing reactive properties, we consider an adaptation of SHML [1]—a syntactic subset of the more expressive μ -calculus logic—describing safety, i.e. monitorable [28], properties. Our choice for this logic was, in part, motivated by the fact that the full μ -calculus had already been adapted to describe concurrent Erlang program behaviour in [22], albeit for model-checking purposes. Given the usual drawbacks associated with full-blown model checking, our work contributes towards an investigation of lightweight verification techniques for μ -calculus properties of Erlang programs. It also allows us to investigate how to extend runtime verification techniques to logics that were not originally intended for this verification setup. More precisely, SHML is a syntactic subset of a larger logic used to specify *branching-time* properties for concurrent systems [2]; our work can thus be used as a first step towards a broader investigation of which other subsets of this branching-time logic can actually be verified at runtime. Crucially, however, SHML acts as an adequate vehicle to study the monitor correctness issues set out in (1) above.

The rest of the paper is structured as follows. Section 2 discusses the formal semantics of our systems and monitor target language. Section 3 discusses reformulations to the logic facilitating the formulation of monitor correctness, discussed later in Sect. 4. Section 5 describes a synthesis algorithm for the logic and a tool built using the algorithm. Subsequently, Sect. 6 proves the correctness of this monitor synthesis. Section 8 concludes.

2 The language

We require a formal semantics for both our monitor-synthesis target language, and the systems we intend to monitor. We partially address this problem by expressing both monitors

Actor Systems, Expressions, Values and Patterns

$A, B, C \in \text{ACTR} ::= i[e \triangleleft q]^m \mid A \parallel B \mid (\nu i)A$	$q, r \in \text{MBox} ::= \epsilon \mid v : q$
$e, d \in \text{EXP} ::= v \mid \text{self} \mid e!d \mid \text{rcv } g \text{ end} \mid e(d) \mid \text{spw } e \mid \text{case } e \text{ of } g \text{ end} \mid x = e, d \mid \dots$	
$v, u \in \text{VAL} ::= x \mid i \mid a \mid \mu y. \lambda x. e \mid \{v, \dots, v\} \mid l \mid \text{exit} \mid \dots$	$l, k \in \text{Lst} ::= \text{nil} \mid v : l$
$p, o \in \text{PAT} ::= x \mid i \mid a \mid \{p, \dots, p\} \mid \text{nil} \mid p : x \mid \dots$	$g, f \in \text{PLst} ::= \epsilon \mid p \rightarrow e; g$

Fig. 1 Erlang syntax

and systems in terms of the same language, i.e. Erlang, thus only requiring one semantics. However, we still need to describe the Erlang tracing semantics we intend to use for our asynchronous monitoring. Although Erlang semantic formalisations exist, e.g. [8, 22, 36], none describe this tracing mechanism. We therefore define a calculus—following [22, 36]—for modelling the *tracing semantics* of a (Turing-complete) subset of the Erlang language (we leave out distribution, process linking and fault-trapping mechanisms).

Figure 1 outlines the language syntax, assuming disjoint denumerable sets of process/actor identifiers $i, j, h \in \text{PID}$, atoms $a, b \in \text{ATOM}$, and variables $x, y, z \in \text{VAR}$. An executing Erlang program is made up of a *system of actors*, ACTR , composed in *parallel*, $A \parallel B$, where some identifiers are local (scoped) to subsystems of actors, and thus not known to the environment, e.g. i in a system $A \parallel (\nu i)B$. Individual actors, denoted as $i[e \triangleleft q]^m$, are uniquely identified by an identifier, i , and consist of an expression, e , executing *wrt.* a local mailbox, q (denoted as a list of values); as we explain later, m denotes the actor monitoring modality. Actor *expressions* typically consist of a sequence of variable binding $x_i = e_i$, terminated by an expression, e_{final} :

$$x_1 = e_1, \dots, x_n = e_n, e_{\text{final}}$$

An expression e_i in a binding $x_i = e_i, e_{i+1}$ is expected to evaluate to a value, v , which is then bound to x_i in the continuation expression e_{i+1} . When instead e_i generates an exception, **exit**, it aborts subsequent computations¹ in e_{i+k} for $1 \leq k \leq (n - i)$. Apart from bindings, expressions may also consist of self references (to the actor’s own identifier), **self**, outputs to other actors, $e_1!e_2$, pattern-matching inputs from the mailbox, **rcv** g **end**, or pattern-matching for case-branches, **case** e **of** g **end** (where g is a list of expressions guarded by patterns, $p_i \rightarrow e_i$), function applications, $e_1(e_2)$, and actor-spawning, **spw** e , amongst others. Values consist of variables, x , process ids, i , recursive functions, $\mu y. \lambda x. e$ (where the preceding μy denotes the binder for function self-reference), tuples $\{v_1, \dots, v_n\}$ and lists, l , amongst others.

Remark 1 The functions $\text{fv}(A)$ and $\text{fId}(A)$ return the free variables and free process identifiers of A *resp.* and are defined in standard fashion. We write $\lambda x. e$ and d, e for $\mu y. \lambda x. e$ and $y = d, e$ *resp.* when $y \notin \text{fv}(e)$. In $p \rightarrow e$, we replace x in p with $_$ whenever $x \notin \text{fv}(e)$. We use standard shorthand for lists of binders and write $\mu y. \lambda(x_1, \dots, x_n). e$ and $(\nu i_1, \dots, i_n) A$ *resp.* for $\mu y. \lambda x_1. \dots \lambda x_n. e$ and $(\nu i_1) \dots (\nu i_n) A$. We sometimes elide mailboxes and write $i[e]$, when these are empty, $i[e \triangleleft \epsilon]$, or when they do not change in the transition rules that follow.

Specific to our formalisation, we also subject each individual actor, $i[e \triangleleft q]^m$, to a *monitoring-modality*, $m, n \in \{\circ, \bullet, *\}$, where \circ, \bullet and $*$ denote *monitored, unmonitored* and *tracing* actors *resp.* Modalities play a crucial role in our language semantics, defined as

¹ Due to **exit** exceptions, variable bindings, $x = e, d$ cannot be encoded as function applications, $\lambda x. d(e)$.

a labelled transition system over systems, $A \xrightarrow{\gamma} B$, where actions $\gamma \in \text{ACT}_\tau$, include bound *output* labels, $(\tilde{h})!v$, and *input* labels, $i?v$ and a distinguished *internal* label, τ . In line with the reactive properties we consider later, our formalisation only traces system interactions with the environment (send and receive messages) from *monitored* actors. Thus, whereas unmonitored, \bullet , and tracing, $*$, actors have standard input and output transition rules

$$\text{SNDU} \frac{m \in \{\bullet, *\}}{j[i!v \triangleleft q]^m \xrightarrow{i!v} j[v \triangleleft q]^m} \quad \text{RCVU} \frac{m \in \{\bullet, *\}}{i[e \triangleleft q]^m \xrightarrow{i?v} i[e \triangleleft q:v]^m}$$

actors with a monitored modality, \circ , i.e. actors j and i in rules SNDM and RCM below, produce a residual message reporting the send and receive interactions ($\{\text{sd}, i, v\}$ and $\{\text{rv}, i, v\}$ resp.) at the tracer’s mailbox i.e. actor h with modality $*$ in the rules below; this models closely the tracing mechanism offered by the Erlang virtual machine (EVM) [11]. In our target language, the list of report messages at the tracer’s mailbox constitutes the system trace to be used for asynchronous monitoring.

$$\text{SNDM} \frac{}{j[i!v \triangleleft q]^\circ \parallel h[d \triangleleft r]^* \xrightarrow{i!v} j[v \triangleleft q]^\circ \parallel h[d \triangleleft r:\{\text{sd}, i, v\}]^*}$$

$$\text{RCVM} \frac{}{i[e \triangleleft q]^\circ \parallel h[d \triangleleft r]^* \xrightarrow{i?v} i[e \triangleleft q:v]^\circ \parallel h[d \triangleleft r:\{\text{rv}, i, v\}]^*}$$

Our LTS semantics assumes *well-formed* actor systems, whereby every actor identifier is *unique*; it is termed to be a *tracing semantics* because a distinguished *tracer* actor, identified by the monitoring modality $*$, receives messages recording external communication events by *monitored* actors. Formally, we write $A \xrightarrow{\gamma} B$ in lieu of $\langle A, \gamma, B \rangle \in \longrightarrow$, the least ternary relation satisfying the rules in Fig. 2. These rules employ *evaluation contexts*, denoted as \mathcal{C} (described below) specifying which sub-expressions are active. For instance, an expression is only evaluated when at the top level variable binding, $x = \mathcal{C}, e$ or when the expression denoting the destination of an output has evaluated to a value, $v!\mathcal{C}$; the other cases are also fairly standard.² We denote the application of a context \mathcal{C} to an expression e as $\mathcal{C}[e]$.

$$\mathcal{C} ::= [-] \mid \mathcal{C}!e \mid v!\mathcal{C} \mid \mathcal{C}(e) \mid v(\mathcal{C}) \mid \text{case } \mathcal{C} \text{ of } g \text{ end} \mid x = \mathcal{C}, e \mid \dots$$

Communication in actor systems happens in two stages: an actor receives messages, keeping them in order in its mailbox, and then selectively reads them at a later stage using pattern-matching—rules RD1 and RD2 describe how mailbox messages are traversed in order to find the first one matching a pattern in the pattern list g , releasing the respective guarded expression e as a result. We choose only to record *external* communication at tracer processes, i.e. between the system and the environment, and do *not* trace internally communication between actors within the system, irrespective of their modality (see COM); structural equivalence rules, $A \equiv B$, are employed to simplify the presentation of our rules—see rule STR and the corresponding structural rules. In PAR, the side-condition enforces the *single-receiver* property, inherent to actor systems; for instance, it prevents a transition with an action $j!v$ when actor j is part of the actor system B . We note that in rule SPW, spawned actors inherit monitorability when launched by a monitored actor, but are launched as unmonitored otherwise.

² In our formalisation, expressions are not allowed to evaluate under a spawn context, $\text{spw}[-]$; this aspect differs from standard Erlang semantics but allows a lightweight description of function application spawning. An adjustment in line with the actual Erlang spawning would be straightforward.

$$\begin{array}{c}
 \text{SNDM} \frac{}{j[C[i!v] \triangleleft q]^\circ \parallel h[d \triangleleft r]^* \xrightarrow{i!v} j[C[v] \triangleleft q]^\circ \parallel h[d \triangleleft (r:(sd, i, v))]^*} \\
 \text{RcvM} \frac{\text{fv}(v) = \emptyset}{i[e \triangleleft q]^\circ \parallel h[d \triangleleft r]^* \xrightarrow{i?v} i[e \triangleleft q:v]^\circ \parallel h[d \triangleleft (r:(rv, i, v))]^*} \\
 \text{SNDU} \frac{m \in \{\bullet, *\}}{j[C[i!v] \triangleleft q]^m \xrightarrow{i!v} j[C[v] \triangleleft q]^m} \qquad \text{RcvU} \frac{m \in \{\bullet, *\} \quad \text{fv}(v) = \emptyset}{i[e \triangleleft q]^m \xrightarrow{i?v} i[e \triangleleft q:v]^m} \\
 \text{SCP} \frac{A \xrightarrow{\gamma} B}{(v \ j)A \xrightarrow{\gamma} (v \ j)B} \quad j \notin (\text{obj}(\gamma) \cup \text{subj}(\gamma)) \qquad \text{OPN} \frac{A \xrightarrow{(\tilde{h})!v} B}{(v \ j)A \xrightarrow{(j\tilde{h})!v} B} \quad i \neq j, j \in \text{subj}((\tilde{h})!v) \\
 \text{COM} \frac{}{j[C[i!v] \triangleleft q]^m \parallel i[e \triangleleft q]^n \xrightarrow{\tau} j[C[v] \triangleleft q]^m \parallel i[e \triangleleft q:v]^n} \\
 \text{PAR} \frac{A \xrightarrow{\gamma} A' \quad \text{obj}(\gamma) \cap \text{fld}(B) = \emptyset}{A \parallel B \xrightarrow{\gamma} A' \parallel B} \qquad \text{RD1} \frac{\text{mtch}(g, v) = e}{i[C[\text{rcv } g \text{ end}] \triangleleft (v : q)]^m \xrightarrow{\tau} i[C[e] \triangleleft q]^m} \\
 \text{RD2} \frac{\text{mtch}(g, v) = \perp}{i[C[\text{rcv } g \text{ end}] \triangleleft (v : q)]^m \xrightarrow{\tau} i[C[e] \triangleleft (v : r)]^m} \\
 \text{CS1} \frac{\text{mtch}(g, v) = e}{i[C[\text{case } v \text{ of } g \text{ end}]]^m \xrightarrow{\tau} i[C[e]]^m} \qquad \text{CS2} \frac{\text{mtch}(g, v) = \perp}{i[C[\text{case } v \text{ of } g \text{ end}]]^m \xrightarrow{\tau} i[C[\text{exit}]]^m} \\
 \text{ASS} \frac{v \neq \text{exit}}{i[C[x = v, e]]^m \xrightarrow{\tau} i[C[e\{v/x\}]]^m} \qquad \text{EXT} \frac{}{i[C[x = \text{exit}, e]]^m \xrightarrow{\tau} i[C[\text{exit}]]^m} \\
 \text{APP} \frac{}{i[C[\mu y. \lambda x. e(v)]]^m \xrightarrow{\tau} i[C[e\{\mu y. \lambda x. e\}_y\{v/x\}]]^m} \qquad \text{SLF} \frac{}{i[C[\text{self}]]^m \xrightarrow{\tau} i[C[i]]^m} \\
 \text{SPW} \frac{(m = \circ = n) \text{ or } (n = \bullet)}{i[C[\text{spw } e] \triangleleft q]^m \xrightarrow{\tau} (v \ j)(i[C[j] \triangleleft q]^m \parallel j[e \triangleleft \epsilon]^m)} \qquad \text{STR} \frac{A \equiv A' \xrightarrow{\gamma} B' \equiv B}{A \xrightarrow{\gamma} B} \\
 \text{sCOM} \frac{}{A \parallel B \equiv B \parallel A} \qquad \text{sASS} \frac{}{(A \parallel B) \parallel C \equiv A \parallel (B \parallel C)} \qquad \text{sCTXP} \frac{A \equiv B}{A \parallel C \equiv B \parallel C} \\
 \text{sEXT} \frac{i \notin \text{fld}(A)}{A \parallel (v \ i)B \equiv (v \ i)(B \parallel A)} \qquad \text{sSWP} \frac{}{(v \ i)(v \ j)A \equiv (v \ j)(v \ i)A} \qquad \text{sCTXS} \frac{A \equiv B}{(vi)A \equiv (vi)B}
 \end{array}$$

Fig. 2 Erlang Semantics for Actor Systems

Mailbox reading—defined by the rules RD1 and RD2 in Fig. 2—includes pattern-matching functionality, allowing the actor to selectively choose which messages to read first from its mailbox whenever the first pattern $p \rightarrow e$ from the pattern list g is matched, returning $e\sigma$, where σ substitutes free variables in e for value binding resulting from the pattern-match; when no pattern is matched, mailbox reading *blocks* - see Definition 1.

Definition 1 (Pattern-Matching) We define $\text{mtch} : \text{PLST} \times \text{VAL} \rightarrow \text{EXP}_\perp$ and $\text{vmtch} : \text{PAT} \times \text{VAL} \rightarrow \text{SUB}_\perp$ as follows:

$$\begin{aligned}
 \text{mtch}(g, v) &\stackrel{\text{def}}{=} \begin{cases} e\sigma & \text{if } g = p \rightarrow e : f, \text{vmtch}(p, v) = \sigma \\ d & \text{if } g = p \rightarrow e : f, \text{vmtch}(p, v) = \perp, \text{mtch}(f, v) = d \\ \perp & \text{otherwise} \end{cases} \\
 \text{vmtch}(p, v) &\stackrel{\text{def}}{=} \begin{cases} \emptyset & \text{if } p = v \text{ (whenever } p \text{ is } a, i \text{ or nil)} \\ \{v/x\} & \text{if } p = x \\ \biguplus_{i=1}^n \sigma_i & \text{if } p = \{p_1, \dots, p_n\}, v = \{v_1, \dots, v_n\} \text{ where} \\ & \text{vmtch}(p_i, v_i) = \sigma_i \\ \sigma \uplus \{l/x\} & \text{if } p = o : x, v = u : l \text{ where } \text{vmtch}(o, u) = \sigma \\ \perp & \text{otherwise} \end{cases} \\
 \sigma_1 \uplus \sigma_2 &\stackrel{\text{def}}{=} \begin{cases} \sigma_1 \cup \sigma_2 & \text{if } \text{dom}(\sigma_1) \cap \text{dom}(\sigma_2) = \emptyset \\ \sigma_1 \cup \sigma_2 & \text{if } \forall v \in \text{dom}(\sigma_1) \cap \text{dom}(\sigma_2). \sigma_1(v) = \sigma_2(v) \\ \perp & \text{if } \sigma_1 = \perp \text{ or } \sigma_2 = \perp \\ \perp & \text{otherwise} \end{cases}
 \end{aligned}$$

Branching for mailbox pattern-matching differs from pattern-match branching for the case construct, described by the rules CS1 and CS2 in Fig. 2: similar to the mailbox read construct, it matches a value to the first appropriate pattern in the pattern list, launching the respective guarded expression with the appropriate variable bindings resulting from the pattern-match; if, however, no match is found it generates an exception, `exit`, which aborts subsequent computation, EXT. The rest of the transition rules, such as APP for function application and SLF for retrieving the actor name, are fairly standard.

Remark 2 Our tracing semantics sits at a higher level of abstraction than that offered by the EVM [11] because trace entries typically contain more information. For instance, the EVM records internal communication between monitored actors, as an output trace entry immediately followed by the corresponding input trace entry; we here describe *sanitised* traces whereby consecutive matching trace entries are filtered out.

Example 1 (Non-deterministic behaviour) Our systems exhibit non-deterministic behaviour through either internal or external choices [25, 30]. Consider the actor system:

$$A \triangleq (v \ j_1, j_2, h)(i[\text{rcv } x \rightarrow \text{obs!}x \text{ end } \triangleleft \epsilon]^\circ \parallel j_1[i!v]^\circ \parallel j_2[i!u]^\circ \parallel h[e \triangleleft q]^*)$$

Actors j_1, j_2 and h are local, i.e. scoped through the construct $(v \ j_1, j_2, h)(\dots)$, thus not visible to the environment. The monitored actor i may receive value v internally from actor j_1 , (2) by rule COM, read it from its mailbox, (3) by RD1, and then output it to some environment actor `obs`, (4) by SNDM, while recording this *external* output at h 's mailbox (the tracer).

$$A \xrightarrow{\tau} (v \ j_1, j_2, h)(i[\text{rcv } x \rightarrow \text{obs!}x \text{ end } \triangleleft v]^\circ \parallel j_1[v]^\circ \parallel j_2[i!u]^\circ \parallel h[e \triangleleft q]^*) \tag{2}$$

$$\xrightarrow{\tau} (v \ j_1, j_2, h)(i[\text{obs!}v \triangleleft \epsilon]^\circ \parallel j_1[v]^\circ \parallel j_2[i!u]^\circ \parallel h[e \triangleleft q]^*) \tag{3}$$

$$\xrightarrow{\text{obs!}v} (v \ j_1, j_2, h)(i[v \triangleleft \epsilon]^\circ \parallel j_1[v] \parallel j_2[i!u] \parallel h[e \triangleleft q : \{\text{sd}, \text{obs}, v\}]^*) \tag{4}$$

But if actor j_2 sends its value to i before j_1 , we observe a different external behaviour

$$\begin{aligned}
 A &\xrightarrow{\tau} . \xrightarrow{\tau} . \xrightarrow{\text{obs!}u} (v \ j_1, j_2, h) \\
 &\quad \times (i[u \triangleleft \epsilon]^\circ \parallel j_1[i!v] \parallel j_2[u] \parallel h[e \triangleleft q : \{\text{sd}, \text{obs}, u\}]^*)
 \end{aligned}$$

i.e. A eventually outputs u instead of v to **obs** (accordingly monitor h would hold the entry $\{\mathbf{sd}, \mathbf{obs}, u\}$ instead); these behaviours amount to an *internal choice*.

External choice results when A receives different external inputs: we can derive

$$A \xrightarrow{i?v_1} B_1 \triangleq (v \ j_1, j_2, h) \times (i[\mathbf{rcv} \ x \rightarrow \mathbf{obs}!x \ \mathbf{end} \ \triangleleft \ v_1]^\circ \parallel j_1[i!v]^\circ \parallel j_2[i!u]^\circ \parallel h[e \ \triangleleft \ q : \{\mathbf{rv}, i, v_1\}]^*)$$

but also $A \xrightarrow{i?v_2} B_2$ for some appropriate B_2 . Subsequently, B_1 can only produce the output $B_1 \xrightarrow{\tau}^* \cdot \xrightarrow{\mathbf{obs}!v_1} B_3$ whereas from B_2 can only produce $B_2 \xrightarrow{\tau}^* \cdot \xrightarrow{\mathbf{obs}!v_2} B_4$. Note that, in the first case, h 's mailbox in B_3 is appended by entries $\{\mathbf{rv}, i, v_1\} : \{\mathbf{sd}, \mathbf{obs}, v_1\}$ whereas, in the second case, h 's mailbox in B_4 is appended by $\{\mathbf{rv}, i, v_2\} : \{\mathbf{sd}, \mathbf{obs}, v_2\}$. \square

Example 2 (Infinite Behaviour) Our systems may exhibit infinite behaviour. Actor A (below) may produce an infinite number of output actions $j!v$ (for $v \neq \mathbf{exit}$) through repeated sequences of function applications, (5) by rule APP, outputs, (6) by rule SNDU and variable assignments, (7) by rule ASS.

$$A \triangleq i[(\mu y. \lambda x. z = j!x, y(z))(v) \ \triangleleft \ q]^\bullet \xrightarrow{\tau} i[z = j!v, (\mu y. \lambda x. z' = j!x, y(z'))(z) \ \triangleleft \ q]^\bullet \tag{5}$$

$$\xrightarrow{j!v} i[z = v, (\mu y. \lambda x. z' = j!x, y(z'))(z) \ \triangleleft \ q]^\bullet \tag{6}$$

$$\xrightarrow{\tau} i[(\mu y. \lambda x. z' = j!x, y(z'))(v) \ \triangleleft \ q]^\bullet = A \tag{7}$$

Systems with infinite behaviour may also transition to a different system with each computational step. System B (below) is a slight modification to A that delegates the output to a newly spawned actor at each iteration. Using a similar sequence of transitions to (5), (6) and (7) together with rule SPW we are able to obtain the computation below, whereby the number of actors grow with each iteration.

$$B \triangleq i[(\mu y. \lambda x. z = \mathbf{spw} \ j!x, y(x))(v) \ \triangleleft \ q]^\bullet \left(\xrightarrow{\tau}\right)^3 \cdot \xrightarrow{j!v} (v \ h)(B \parallel h[v \ \triangleleft \ \epsilon]^\bullet) \left(\xrightarrow{\tau}\right)^3 \cdot \xrightarrow{j!v} (v \ h, h')(B \parallel h[v \ \triangleleft \ \epsilon]^\bullet \parallel h'[v \ \triangleleft \ \epsilon]^\bullet) \left(\xrightarrow{\tau}\right)^3 \cdot \xrightarrow{j!v} \dots$$

Such infinite behaviour together with the non-deterministic behaviour discussed in Example 1 typically lead to state-explosion problems when systems are verified exhaustively. \square

3 The logic

In order to specify reactive properties for the actor systems discussed in Sect. 2, we consider an adaptation of SafeHML [1] (sHML), a sub-logic of the Hennessy–Milner Logic (HML) with recursion—the latter logic has been shown to be a reformulation of the expressive μ -calculus [27]. It assumes a denumerable set of formula variables, $X, Y \in \text{LVAR}$, and is inductively defined by the following grammar:

$$\varphi, \psi \in \text{sHML} ::= \mathbf{ff} \mid \varphi \wedge \psi \mid [\alpha]\varphi \mid X \mid \mathbf{max}(X, \varphi)$$

The formulas for falsity, \mathbf{ff} , conjunction, $\varphi \wedge \psi$, and action necessity, $[\alpha]\varphi$, are inherited from HML[25], whereas variables X and the recursion construct $\mathbf{max}(X, \varphi)$ are used to define

maximal fixpoints; as expected, $\max(X, \varphi)$ is a binder for the free variables X in φ , inducing standard notions of open and closed formulas. We only depart from the logic of [1] by limiting formulas to *basic actions* $\alpha, \beta \in \text{BACT}$, including just input, $i?v$, and *unbound* outputs, $i!v$, so as to keep our technical development manageable.

Remark 3 The handling of *bound* output actions, $(\tilde{h})i!v$, is well understood [31] and does not pose problems to monitoring, apart from making action pattern-matching cumbersome (consult [24] for an example of how bound values can be matched); it also complicates instrumentation (see Sects. 4 and 5). Certain (silent) τ labels can also be monitored using minor adaptations (see, for instance, Remark 2); they however increase substantially the size of the traces recorded, unnecessarily cluttering the tracing semantics of Sect. 2.

The semantics of our logic is defined for closed formulas, using the operation $\varphi\{\psi/X\}$, which substitutes free occurrences of X in φ with ψ without introducing any variable capture. It is specified as the satisfaction relation of Definition 2 (adapted from [1]). In what follows, we write *weak* transitions $A \Longrightarrow B$ and $A \xrightarrow{\alpha} B$, for $A \xrightarrow{\tau}^* B$ and $A \xrightarrow{\tau}^* \cdot \xrightarrow{\alpha} \cdot \xrightarrow{\tau}^* B$ resp. We let $s, t \in (\text{BACT})^*$ range over *lists of basic actions* and write sequences of weak actions $A \xrightarrow{\alpha_1} \cdots \xrightarrow{\alpha_n} B$, where $s = \alpha_1, \dots, \alpha_n$, as $A \xrightarrow{s} B$ (or as $A \xrightarrow{s}$ when B is unimportant).

Definition 2 (*Satisfiability*) A relation $\mathcal{R} \in \text{ACTR} \times \text{SHML}$ is a satisfaction relation iff:

$$\begin{aligned} (A, \text{ff}) \in \mathcal{R} & \text{ never} \\ (A, \varphi \wedge \psi) \in \mathcal{R} & \text{ implies } (A, \varphi) \in \mathcal{R} \text{ and } (A, \psi) \in \mathcal{R} \\ (A, [\alpha]\varphi) \in \mathcal{R} & \text{ implies } (B, \varphi) \in \mathcal{R} \text{ whenever } A \xrightarrow{\alpha} B \\ (A, \max(X, \varphi)) \in \mathcal{R} & \text{ implies } (A, \varphi\{\max(X, \varphi)/X\}) \in \mathcal{R} \end{aligned}$$

Satisfiability, \models_s , is the *largest* satisfaction relation; we write $A \models_s \varphi$ for $(A, \varphi) \in \models_s$. It follows from standard fixed-point theory that the implications of satisfaction relation are bi-implications for Satisfiability.

Example 3 (*Satisfiability*) Consider the safety formula

$$\varphi_{\text{safe}} \triangleq \max(X, [\alpha][\alpha][\beta]\text{ff} \wedge [\alpha]X) \tag{8}$$

stating that a satisfying actor system should never perform a sequence of two external actions α followed by the external action β (through the subformula $[\alpha][\alpha][\beta]\text{ff}$), and that this needs to hold after every α action (through $[\alpha]X$); effectively the formula states that trace sequences of α -actions *greater than two* cannot be followed by a β -action.

A system A_1 exhibiting *just* the external behaviour $A_1 \xrightarrow{\alpha\beta}$ satisfies φ_{safe} , as would a system A_2 with just the *infinite* behaviour $A_2 \xrightarrow{\alpha} A_2$. System A_3 , with a trace $A_3 \xrightarrow{\alpha\alpha\beta}$, *does not* satisfy this property, $A_3 \not\models_s \varphi_{\text{safe}}$, according to Definition 2. However, this system may be capable of producing other external traces at runtime (e.g. as a result of some internal choice). In fact, if A_3 exhibits the alternate behaviour $A_3 \xrightarrow{\beta}$, one would be *unable* to observe any violation from A_3 at runtime since the external trace β is allowed by φ_{safe} .

Since actors may violate a property along one execution but satisfy it along another, the inverse of \models_s , i.e. $A \not\models_s \varphi$, is too coarse to be used for a definition of monitor correctness along the lines of (1) discussed earlier. We thus define a *violation relation*, Definition 3, characterising actors violating a property along a specific execution trace.

Definition 3 (*Violation*) The violation relation, denoted as \models_v , is the *least* relation of the form $(\text{ACTR} \times \text{BACT}^* \times \text{SHML})$ satisfying the following rules:

$$\begin{aligned}
 A, s &\models_v \text{ff} \quad \text{always} \\
 A, s &\models_v \varphi \wedge \psi \quad \text{if } A, s \models_v \varphi \text{ or } A, s \models_v \psi \\
 A, \alpha s &\models_v [\alpha]\varphi \quad \text{if } A \xrightarrow{\alpha} B \text{ and } B, s \models_v \varphi \\
 A, s &\models_v \max(X, \varphi) \quad \text{if } A, s \models_v \varphi \{ \max(X, \varphi) / X \}
 \end{aligned}$$

We write $A, s \models_v \varphi$ in lieu of $(A, s, \varphi) \in \models_v$. It also follows from standard fixed-point theory that the constraints of the violation relation are bi-implications.

Example 4 (*Violation*) Recall the safety formula φ_{safe} defined in (8). Actor A_3 , from Example 3, together with the witness violating trace $\alpha\alpha\beta$ violate φ_{safe} , i.e. $(A_3, \alpha\alpha\beta) \models_v \varphi_{\text{safe}}$. However, A_3 together with trace β do not violate φ_{safe} , i.e. $(A_3, \beta) \not\models_v \varphi_{\text{safe}}$. Definition 3 relates a violating trace with an actor *only when* that trace leads the actor to a violation: if A_3 cannot perform the trace $\alpha\alpha\alpha\beta$, by Definition 3, we have $(A_3, \alpha\alpha\alpha\beta) \not\models_v \varphi_{\text{safe}}$, even though the trace is prohibited by φ_{safe} . A violating trace may also lead a system to a violation before its end, e.g. $(A_3, \alpha\alpha\beta\alpha) \models_v \varphi_{\text{safe}}$ according to Definition 3. \square

Despite the technical discrepancies between Definition 2 and Definition 3—e.g. the use of maximal versus minimal fixpoints and a differing model—we show that Definition 3 corresponds, in some sense, to the dual of Definition 2.

Theorem 1 (*Correspondence*) $\exists s.(A, s) \models_v \varphi$ iff $A \not\models_s \varphi$

Proof For the *if* case we prove the contrapositive, i.e. that $\forall s.A, s \not\models_v \varphi$ implies $A \models_s \varphi$ by co-inductively showing that $\mathcal{R} = \{(A, \varphi) \mid \forall s.A, s \not\models_v \varphi\}$ is a satisfaction relation. The proof is by induction on the structure of φ . We outline two cases (see [19] for the rest):

- $\varphi \wedge \psi$: From the definition of \mathcal{R} we know that $\forall s.A, s \not\models_v \varphi \wedge \psi$ and, by Definition 3, this implies that $\forall s.(A, s \not\models_v \varphi \text{ and } A, s \not\models_v \psi)$. Distributing the universal quantification yields $\forall s.A, s \not\models_v \varphi$ and $\forall s.A, s \not\models_v \psi$, and by the definition of \mathcal{R} we obtain $(A, \varphi) \in \mathcal{R}$ and $(A, \psi) \in \mathcal{R}$, as required for satisfiability relations by Definition 2.
- $[\alpha]\varphi$: From the definition of \mathcal{R} we know that $\forall s.A, s \not\models_v [\alpha]\varphi$. In particular, for all $s = \alpha t$, we know that $\forall t.A, \alpha t \not\models_v [\alpha]\varphi$. From Definition 3 it must be the case that whenever $A \xrightarrow{\alpha} B$ we have that $\forall t.B, t \not\models_v \varphi$, which in turn implies that $(B, \varphi) \in \mathcal{R}$ (from the definition of \mathcal{R}); this is the implication required by Definition 2.

For the *only-if* case we prove $\exists s.A, s \models_v \varphi$ implies $A \not\models_s \varphi$ by rule induction on $A, s \models_v \varphi$. Note that $A \not\models_s \varphi$ means that there does not exist *any* satisfiability relation including the pair (A, φ) . Again we outline the main cases (see [19] for the rest):

- $A, s \models_v [\alpha]\varphi$ because $s = \alpha s', A \xrightarrow{\alpha} B$ and $B, s' \models_v \varphi$: By $B, s' \models_v \varphi$ and I.H. we obtain $B \not\models_s \varphi$, and subsequently, by $A \xrightarrow{\alpha} B$, we conclude that $A \not\models_s [\alpha]\varphi$.
- $A, s \models_v \max(X, \varphi)$ because $A, s \models_v \varphi \{ \max(X, \varphi) / X \}$: By $A, s \models_v \varphi \{ \max(X, \varphi) / X \}$ and I.H. we obtain $A \not\models_s \varphi \{ \max(X, \varphi) / X \}$ which, in turn, implies that $A \not\models_s \max(X, \varphi)$. \square

Definition 3 and Theorem 1 allows us to show that *every* formula $\varphi \in \text{SHML}$ denotes a *safety language*, as defined in [12,28].³ Informally, this states that whenever we determine

³ Note that we do *not* show that SHML captures all the safety properties expressible in HML with recursion, and there are in fact other formulas that specify safety properties such as tt.

that $A \not\models_s \varphi$ along a particular trace s , such a judgement is preserved for all extensions of that trace. This property, in turn, implies that all formulas in sHML are monitorable [5]. We formally prove this result in Theorem 2, assuming the standard notion of trace prefixes i.e. $s \leq t$ iff $\exists s'$ such that $t = ss'$.

Theorem 2 (Safety Properties and sHML) $A, s \models_v \varphi$ and $s \leq t$ implies $A, t \models_v \varphi$

Proof By rule induction on $A, s \models_v \varphi$. We outline the main cases leaving the rest for [19]:

$A, s \models_v \varphi \wedge \psi$ because $A, s \models_v \varphi$: By $A, s \models_v \varphi, s \leq t$ and I.H. we obtain $A, t \models_v \varphi$ which, by the same rule, implies $A, t \models_v \varphi \wedge \psi$.

$A, s \models_v [\alpha]\varphi$ because $s = \alpha s', A \xrightarrow{\alpha} B$ and $B, s' \models_v \varphi$: From $s = \alpha s'$ and $s \leq t$ we know that $s' \leq t'$ for some t' where $t = \alpha t'$. By $s' \leq t', B, s' \models_v \varphi$ and I.H. we obtain $B, t' \models_v \varphi$ and by $A \xrightarrow{\alpha} B$ and $t = \alpha t'$ we derive $A, t \models_v [\alpha]\varphi$. \square

4 Correctness

Specifying online monitor correctness is complicated by the fact that, in general, we have limited control over the behaviour of the systems being monitored. For starters, a system that does not satisfy a property may still exhibit runtime behaviour that does not violate it, as discussed earlier in the case of system A_3 of Example 3 and Example 4. We deal with system non-determinism by only requiring monitor detection when the system performs a violating execution: this can be expressed through the violation relation of Definition 3.

At runtime, a system may also interfere with the execution of monitors. Appropriate instrumentation can limit system effects on the monitors. In our asynchronous actor setting, direct interferences from the system to the monitors can be precluded by (i) locating the monitors at process identifiers not known to the system (ii) preventing the monitors from communicating these identifiers to the system. These measures inhibit the system’s ability to send messages to the monitors.

A system may also interfere with monitor executions indirectly by diverging, i.e. infinite internal computation (τ -transitions) without external actions. This can prevent the monitors from progressing during their execution and thus postpone indefinitely violation detections [32]. In our case, divergence is handled, in part, by the EVM itself, which guarantees fair executions for concurrent actors [11]. In settings where fair executions may be assumed, it suffices to require a weaker property for monitors, reminiscent of the condition in fair/should-testing [33]. Definition 4 states that, for an arbitrary basic action α , an actor system A satisfies the predicate *should- α* if, for all sequences of internal actions leading to some system B , there always exists an execution from B that can produce the action α ; in the case of monitors, the external should-action is set to a reserved violation-detection action, e.g., fail!

Definition 4 (Should- α) $A \Downarrow_\alpha \stackrel{\text{def}}{=} (A \Longrightarrow B \text{ implies } B \xrightarrow{\alpha})$

We limit monitoring to monitorable systems, where all actors are subject to a monitorable modality, i.e. modality \circ .

$$A \equiv (v \tilde{h})(i[e \triangleleft q]^m \parallel B) \text{ implies } m = \circ$$

This guarantees that (i) they can be composed with a tracer actor (ii) all the basic actions produced by the system are recorded as trace entries at the tracer’s mailbox.⁴ Monitor correctness

⁴ Due to asynchronous communication, even scoped actors can produce visible actions by sending messages to environment actors.

is defined for (unmonitored) *basic* systems, satisfying the condition:

$$A \equiv (\nu \tilde{h})(i[e \triangleleft q]^m \parallel B) \text{ implies } m = \bullet$$

which are instrumented to execute in parallel with the monitor. Our instrumentation is defined through the operation $\lceil - \rceil$, Definition 5, converting basic systems to monitorable ones using `trace/2` and `set_on_spawn` Erlang commands [11]; see Lemma 1. Importantly, instrumentation does not affect the visible behaviour of a basic system; see Lemma 2.

Definition 5 (*Instrumentation*) $\lceil - \rceil : \text{ACTR} \rightarrow \text{ACTR}$ is defined inductively as:

$$\lceil i[e \triangleleft q]^m \rceil \stackrel{\text{def}}{=} i[e \triangleleft q]^\circ \quad \lceil B \parallel C \rceil \stackrel{\text{def}}{=} \lceil B \rceil \parallel \lceil C \rceil \quad \lceil (\nu i)B \rceil \stackrel{\text{def}}{=} (\nu i)\lceil B \rceil$$

Lemma 1 *If A is a basic system then $\lceil A \rceil$ is monitorable.*

Lemma 2 *For all basic actors A where $i \notin \text{fld}(A)$:*

$$A \xrightarrow{\alpha} B \text{ iff } \begin{cases} (\nu i)(\lceil A \rceil \parallel i[e \triangleleft q]^*) \xrightarrow{j!v} (\nu i)(\lceil B \rceil \parallel i[e \triangleleft q : \{\text{sd}, j, v\}]^*) & \text{if } \alpha = j!v \\ (\nu i)(\lceil A \rceil \parallel i[e \triangleleft q]^*) \xrightarrow{j?v} (\nu i)(\lceil B \rceil \parallel i[e \triangleleft q : \{\text{rv}, j, v\}]^*) & \text{if } \alpha = j?v \\ (\nu i)(\lceil A \rceil \parallel i[e \triangleleft q]^*) \xrightarrow{\tau} (\nu i)(\lceil B \rceil \parallel i[e \triangleleft q]^*) & \text{if } \alpha = \tau \end{cases}$$

We are now in a position to state monitor correctness, for some predefined violation-detection monitor action `fail!`, Definition 6; in what follows, `fail` is always assumed to be fresh. We restrict our definition to expressions e located at a fresh scoped location i (not used by the system, i.e. $i \notin \text{fld}(A)$) with an empty mailbox, ϵ ; expression e may then spawn concurrent submonitors while executing. The definition can be extended to generic concurrent monitors, i.e. multiple expressions, in straightforward fashion.

Definition 6 (*Correctness*) $e \in \text{EXP}$ is a correct monitor for $\varphi \in \text{SHML}$ iff for any basic actors $A \in \text{ACTR}$, $i \notin \text{fld}(A)$, and execution traces $s \in (\text{ACT} \setminus \{\text{fail!}\})^*$:

$$(\nu i)(\lceil A \rceil \parallel i[e]^*) \xrightarrow{s} B \text{ implies } (A, s \models_v \varphi \text{ iff } B \Downarrow_{\text{fail!}})$$

Definition 6 states that e correctly monitors property φ whenever, for any trace of environment interactions, s , of a monitored system, $(\nu i)(\lceil A \rceil \parallel i[e \triangleleft \epsilon]^*)$, yielding system B , if s leads A to a violation of φ , then system B should always detect it, and viceversa. It formalises the definition of monitor correctness outlined earlier in (1) from Sect. 1.

5 Automated monitor synthesis

We define a translation from SHML formulas to Erlang *monitors* that asynchronously analyse a system and flag an alert whenever they detect violations by the current system execution (for the respective SHML formula). This translation describes the core algorithm for a tool automating monitor synthesis from SHML formulas [21].

Despite its relative simplicity, SHML still provides opportunities for performing concurrent monitoring. The most obvious case is the translation of conjunction formulas, $\varphi_1 \wedge \varphi_2$, whereby the resulting code needs to check *both* sub-formulas φ_1 and φ_2 so as to ensure that *neither* is violated; since conjunctions are prevalent in many monitoring logics, we conjecture that the concepts discussed here extend in straightforward fashion to similar runtime verification settings with asynchronous concurrent monitors. More specifically, a translation

in terms of two *concurrent* (sub)monitors, each analysing different parts of the trace so as to ensure the observation of its respective sub-formula, constitutes a natural synthesis of the conjunction formula in our target language: it adheres to recommended Erlang practices advocating for concurrency wherever possible [11], but also allows us to benefit from the advantages of concurrent monitors discussed in Sect. 1.

Example 5 (Conjunction Formulas) Consider the two sHML formulas

$$\begin{aligned} \varphi_{\text{no_dup_ans}} &\triangleq [\alpha_{\text{call}}] (\max(X, [\beta_{\text{ans}}][\beta_{\text{ans}}] \text{ff} \wedge [\beta_{\text{ans}}][\alpha_{\text{call}}] X)) \\ \varphi_{\text{react_ans}} &\triangleq \max(Y, [\beta_{\text{ans}}] \text{ff} \wedge [\alpha_{\text{call}}][\beta_{\text{ans}}] Y) \end{aligned}$$

Formula $\varphi_{\text{no_dup_ans}}$ requires that call actions α_{call} are at most serviced by a *single* answer action β_{ans} , and that this condition is invariant for any sequence of $(\alpha_{\text{call}}, \beta_{\text{ans}})$ pairs. On the other hand, formula $\varphi_{\text{react_ans}}$ requires that answer actions are only produced *in response to* call actions; again this is required to be invariant for sequences of $(\alpha_{\text{call}}, \beta_{\text{ans}})$ pairs. Although one can rephrase the conjunction of the two formulas as a formula without a top-level conjunction, it is more straightforward to use two concurrent monitors executing in parallel (one for each sub-formula in $\varphi_{\text{no_dup_ans}} \wedge \varphi_{\text{react_ans}}$). There are also other reasons why it would be beneficial to keep the sub-formulas separate: for instance, keeping the formulas disentangled improves maintainability and separation of concerns when subformulas originate from distinct requirement specifications.⁵ □

Multiple conjunctions also arise indirectly when used under fix-point operators. When synthesising concurrent monitors analysing separate branches of such recursive properties, it is important to generate monitors that can dynamic spawn further submonitors themselves as required at runtime, so as to keep the monitoring overheads to a minimum.

Example 6 (Conjunctions and Fixpoints) Recall φ_{safe} , from (8) in Example 3.

$$\varphi_{\text{safe}} \triangleq \max(X, [\alpha][\alpha][\beta] \text{ff} \wedge [\alpha] X)$$

Semantically, the formula represents the infinite-depth tree with an infinite number of conjunctions, as depicted in Fig. 3a. Although in practice, we cannot generate an infinite number of concurrent monitors, φ_{safe} will translate into possibly more than two concurrent monitors executing in parallel. The same applies for the conjunctions used in formulas $\varphi_{\text{no_dup_ans}}$ and $\varphi_{\text{react_ans}}$ from Example 5. □

Our monitor synthesis, $\llbracket - \rrbracket^{\text{m}} : \text{sHML} \rightarrow \text{EXP}$, given in Definition 7, takes a *closed, guarded*⁶ sHML formula and returns an Erlang function. This function takes a *map* as an argument (encoded as a list of tuples from formula variables to other synthesised monitors of the same form) and releases an expression that performs the monitoring.

The map encodes an *environment* that maps formula variables to logical formula, introduced by the binding in the construct $\max(X, \varphi)$; it is used for *lazy* recursive unrolling of formulas so as to minimize monitoring overhead. For instance, when synthesising formula φ_{safe} from Example 3, the algorithm initially spawns only two concurrent submonitors, one checking for the subformula $[\alpha][\alpha][\beta] \text{ff}$, and another one checking for the formula $[\alpha] X$, as is depicted in Fig. 3b. Whenever the rightmost submonitor in Fig. 3b observes the action α and reaches X , it consults its environment and retrieves the respective formula bound to

⁵ One potential disadvantage of splitting formulas is that of increasing communication amongst monitors.

⁶ In *guarded* sHML formulas, variables appear only as a sub-formula of a necessity formula.

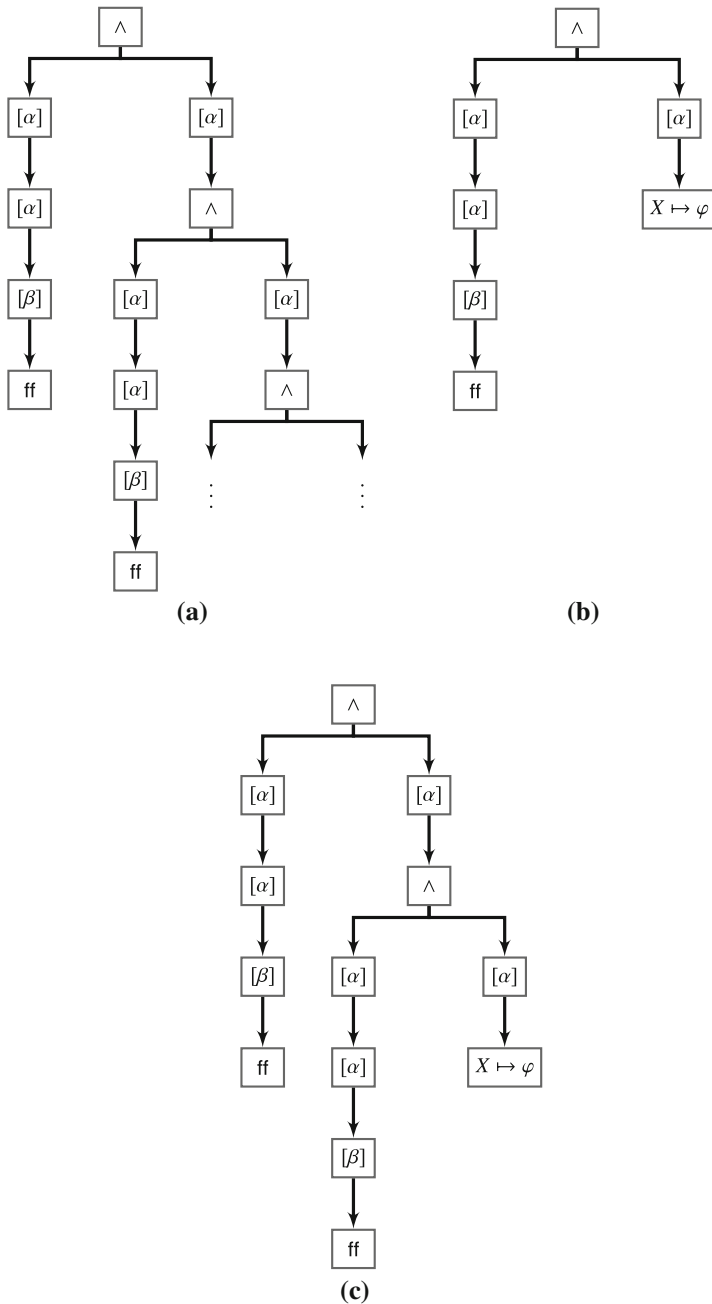


Fig. 3 Monitor Combinator generation for φ_{safe} of Ex. 3. **a** Denotation of φ_{safe} defined in (8). **b** Constructed concurrent monitors for φ_{safe} where $\varphi = [\alpha][\alpha][\beta]\text{ff}\wedge[\alpha]X$. **c** First expansion of the constructed monitor for φ_{safe}

X ; this allows it to unfold X and spawns an additional submonitor as depicted in Fig. 3c, thereby increasing monitor overheads *incrementally* as needed.

As discussed earlier in Sect. 2, traces are encoded as messages ordered inside the respective mailbox of the monitor. Monitoring thus involves reading these messages from the mailbox *in order*, analysing them, and determining what action to take. In a system of hierarchically-organised concurrent monitors, monitor actions can either flag a violation, forward messages to other sub-monitors, or spawn new monitors.

Definition 7 (*Synthesis*) $\llbracket - \rrbracket^m$ is defined on the structure of the SHML formula:

$$\begin{aligned} \llbracket \text{ff} \rrbracket^m &\stackrel{\text{def}}{=} \lambda x_{\text{env}}. \text{fail!} \\ \llbracket \varphi_1 \wedge \varphi_2 \rrbracket^m &\stackrel{\text{def}}{=} \begin{cases} \lambda x_{\text{env}}. y_{\text{pid}1} = \text{spw} (\llbracket \varphi_1 \rrbracket^m(x_{\text{env}})), \\ y_{\text{pid}2} = \text{spw} (\llbracket \varphi_2 \rrbracket^m(x_{\text{env}})), \\ \text{fork}(y_{\text{pid}1}, y_{\text{pid}2}) \end{cases} \\ \llbracket [\alpha] \varphi \rrbracket^m &\stackrel{\text{def}}{=} \begin{cases} \lambda x_{\text{env}}. \text{rcv} \\ \quad \text{tr}(\alpha) \rightarrow \llbracket \varphi \rrbracket^m(x_{\text{env}}); \\ \quad _ \rightarrow \text{stop} \\ \text{end} \end{cases} \\ \llbracket \max(X, \varphi) \rrbracket^m &\stackrel{\text{def}}{=} \lambda x_{\text{env}}. y_{\text{mon}} = \llbracket \varphi \rrbracket^m, y_{\text{mon}}(\{X', y_{\text{mon}}\} : x_{\text{env}}) \\ \llbracket X \rrbracket^m &\stackrel{\text{def}}{=} \lambda x_{\text{env}}. y_{\text{mon}} = \text{lookUp}(X', x_{\text{env}}, y_{\text{mon}}(x_{\text{env}})) \end{aligned}$$

Auxiliary function definitions and meta-operators:

$$\begin{aligned} \text{fork} &\stackrel{\text{def}}{=} \mu y_{\text{rec}}. \lambda (x_{\text{pid}1}, x_{\text{pid}2}). \text{rcv } z \rightarrow (x_{\text{pid}1}!z, x_{\text{pid}2}!z) \text{ end}, y_{\text{rec}}(x_{\text{pid}1}, x_{\text{pid}2}) \\ \text{lookUp} &\stackrel{\text{def}}{=} \begin{cases} \mu y_{\text{rec}}. \lambda (x_{\text{var}}, x_{\text{map}}). \text{case } x_{\text{map}} \text{ of } \{ \{x_{\text{var}}, z_{\text{mon}}\} : _ \} \rightarrow z_{\text{mon}} \\ \quad _ : z_{\text{tl}} \rightarrow y_{\text{rec}}(x_{\text{var}}, z_{\text{tl}}) \\ \quad \text{nil} \rightarrow \text{exit} \\ \text{end} \end{cases} \end{aligned}$$

In Definition 7, the synthesised monitor for formula ff immediately reports a violation to some supervisor actor handling the violation; we assume that the supervisor actor is identified by the name fail . Conjunction formulas, $\varphi_1 \wedge \varphi_2$, translate into the spawning of the respective monitors for φ_1 and φ_2 , i.e. the command $y_{\text{pid}i} = \text{spw} (\llbracket \varphi_i \rrbracket^m(x_{\text{env}}))$ in Definition 7, and the subsequent forwarding of trace messages to these spawned monitors through the auxiliary function fork . The translated monitor for $[\alpha]\varphi$ behaves as the monitor translation for φ once it receives a trace message encoding the occurrence of action α , i.e. the guarded expression $\text{tr}(\alpha) \rightarrow \llbracket \varphi \rrbracket^m(x_{\text{env}})$ in Definition 7, which uses the meta-function $\text{tr}(_)$ defined below:

$$\text{tr}(i?v) \stackrel{\text{def}}{=} \{rv, i, v\} \quad \text{tr}(i!v) \stackrel{\text{def}}{=} \{sd, i, v\}$$

Importantly, the monitor for $[\alpha]\varphi$ terminates if the trace message does not correspond to α , i.e. the guarded expression $_ \rightarrow \text{stop}$ in Definition 7.

The translations for formulas $\max(X, \varphi)$ and X are best understood together. The monitor for $\max(X, \varphi)$ behaves like that for φ , under the *extended* environment where X is mapped to the monitor for φ , i.e. the operation $\{X', y_{\text{mon}}\} : x_{\text{env}}$ in Definition 7; this effectively models the formula unrolling $\varphi\{\max(X, \varphi)/X\}$ from Definition 3. The monitor for X retrieves the respective monitor translation bound to X in the map using the auxiliary function lookUp ,

$$y_{\text{mon}}(\{Y', y_{\text{mon}}\} : x_{\text{env}})$$

□

5.1 Tool implementation

We have constructed a tool called **detectEr** [21], that implements the monitor synthesis of Definition 7: given an SHML formula it generates a monitor that can be instrumented with minimal changes to the system (actual Erlang code), as discussed earlier in Sect. 4. Despite not being the main focus of this work, we conducted preliminary empirical experiments evaluating the performance of our synthesised monitors. This was carried out using a simulated server that launches individual workers to handle a series of requests from individual clients; we also injected faults making certain workers non-deterministically behave erratically. We synthesised monitors to check that each worker respects the *no-duplicate-reply* property from Example 5:

$$\varphi_{\text{wrkr}} \triangleq [\text{wrk?req}] (\max(X, [\text{clnt!rply}] [\text{clnt!rply}] \text{ff} \wedge [\text{clnt!rply}] [\text{wrk?req}] X)$$

and calculated the overheads incurred for varying number of client requests (i.e. concurrent workers); we also compared this with the performance a monitor that checks for property violations in sequential fashion. Tests were carried out on an Intel Core i7 processor with 8GB of RAM, running Microsoft Windows 8 and EVM version R15B02. The results, summarised in the table below, show that our synthesised concurrent monitoring yields acceptable overheads that are consistently lower than those of a sequential monitor. We conjecture that this discrepancy can be increased further when monitoring for recursive properties with longer chains of necessity formulas. For more extensive empirical results relating to the tool **detectEr**, see [9].

No. of reqs.	Sequential			Concurrent		
	Time (μs)	Time (μs)	Ovhd (%)	Time (μs)	Ovhd. (%)	Improv. (%)
250	117.813	121.667	3.27	118.293	0.40	2.86
350	185.232	202.500	9.32	194.793	5.16	4.16
450	237.606	248.333	4.51	242.380	2.01	2.51
550	286.461	319.167	11.42	308.853	7.82	3.60
650	345.543	372.232	7.72	354.333	2.54	5.18

6 Proving correctness

The preliminary results obtained in Sect. 5 advocate for the feasibility of using concurrent monitors. We however still need to show that the monitors synthesised are correct. Definition 6 allows us to state one of the main results of the paper, Theorem 3.

Theorem 3 (Correctness) *For all $\varphi \in \text{SHML}$, $\text{Mon}(\varphi)$ is a correct monitor for φ .*

Proving Theorem 3 directly can be an arduous task: for *any* SHML formula, it requires reasoning about *all* the possible execution paths of *any* monitored system in parallel with the instrumented monitor. We propose a formal technique for alleviating the task of ascertaining

the monitor correctness of Definition 6 by teasing apart *three* separate (weaker) monitor-conditions: they are referred to as *Violation Detectability*, *Detection Preservation* and *Monitor Separability*.

These conditions are important properties in their own right—for instance, detection preservation requires the monitor to behave *deterministically* wrt. violation detections, whereas Monitor Separability requires that the monitor computation *does not affect* the execution of the monitored system. Moreover, the three conditions pose advantages to the checking of monitor correctness: since these conditions are independent to one another, they can be checked in parallel by distinct analysing entities; alternatively, the conditions that are inexpensive to check may be carried out before the more expensive ones, thus acting as vetting phases that abort early and keep the analysis cost to a minimum. More importantly though, the three conditions together imply our original monitor-correctness criteria.

The first sub-property is *Violation Detectability*, Lemma 3, guaranteeing that every violating trace s of formula φ is detectable by the respective synthesised monitor,⁷ (the *only-if* case) and that there are no false detections (the *if* case). This property is easier to verify than Theorem 3 since it requires us to consider the execution of the monitor in isolation and, more importantly, requires us to verify the *existence of an execution path* that detects the violation; concurrent monitors typically have multiple execution paths and Theorem 3 requires us to prove this property for *all* of the possible monitor execution paths.

Lemma 3 (Violation Detectability) *For basic $A \in \text{ACTR}$ and $i \notin \text{fld}(A)$, $A \xrightarrow{s}$ implies:*

$$A, s \models_v \varphi \text{ iff } i[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \xrightarrow{\text{fail!}}$$

Detection preservation (Lemma 4), the second sub-property, is not concerned with relating detections to the violations specified by our logic semantics, $A, s \models_v \varphi$. Instead it guarantees that if a monitor can potentially detect a violation, further reductions do not exclude the possibility of this detection. In the case where monitors always have a finite reduction wrt. their mailbox contents (as it turns out to be the case for monitors synthesised by Definition 7) this condition guarantees that the monitor will detect violations *deterministically*. More generally, however, in a setting that guarantees fair actor executions, Lemma 4 ensures that detection will always eventually occur, even when monitors execute in parallel with other, potentially divergent, systems.

Lemma 4 (Detection Preservation) *For all $\varphi \in \text{SHML}$, $q \in \text{VAL}^*$*

$$\left(i[\text{Mon}(\varphi) \triangleleft q]^* \xrightarrow{\text{fail!}} \text{ and } i[\text{Mon}(\varphi) \triangleleft q]^* \Longrightarrow B \right) \text{ implies } B \xrightarrow{\text{fail!}}$$

The third sub-property is Separability, Lemma 5, which implies that the behaviour of a (monitored) system *is independent of* the monitor and, dually, the behaviour of the monitor depends, *at most*, on the trace generated by the system.

Lemma 5 (Monitor Separability) *For all basic $A \in \text{ACTR}$, $i \notin \text{fld}(A)$, $\varphi \in \text{SHML}$, and $s \in (\text{ACT} \setminus \{\text{fail!}\})^*$,*

$$(\forall i) (\lceil A \rceil \parallel i[\text{Mon}(\varphi)]^*) \xrightarrow{s} B \text{ implies } \exists B', B'' \text{ s.t.}$$

$$B \equiv (\forall i) (B' \parallel B'') \text{ and } A \xrightarrow{s} A' \text{ s.t. } B' = \lceil A' \rceil \text{ and } i[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \Longrightarrow B''$$

⁷ We elevate tr to basic action sequences s in pointwise fashion, $\text{tr}(s)$, where $\text{tr}(\epsilon) = \epsilon$.

These three properties suffice to show monitor correctness.

Theorem 3 (Correctness). *For all $\varphi \in \text{sHML}$, $\text{Mon}(\varphi)$ is a correct monitor for φ .*

Proof According to Definition 6 we have to show:

$$(\nu i)(\lceil A \rceil \parallel i[\text{Mon}(\varphi)]^*) \xrightarrow{s} B \text{ implies } (A, s \models_{\nu} \varphi \text{ iff } B \Downarrow_{\text{fail!}})$$

We assume that $(\nu i)(\lceil A \rceil \parallel i[\text{Mon}(\varphi)]^*) \xrightarrow{s} B$ holds and consider the two sides of the bi-implication separately. For the *only-if* case, we assume

$$(\nu i)(A \parallel i[\text{Mon}(\varphi)]^*) \xrightarrow{s} B \tag{9}$$

$$A, s \models_{\nu} \varphi \tag{10}$$

In order to show $B \Downarrow_{\text{fail!}}$, we use Definition 4 (*Should- α*) to expand $B \Downarrow_{\text{fail!}}$. We thus also assume $B \Longrightarrow B'$, for arbitrary B' , and then be required to prove that $B' \xrightarrow{\text{fail!}}$. From (9), $B \Longrightarrow B'$ and Lemma 5 (Monitor Separability) we know

$$\exists B'', B''' \text{ s.t. } B' \equiv (\nu i)(B'' \parallel B''') \tag{11}$$

$$A \xrightarrow{s} A' \text{ for some } A' \text{ where } \lceil A' \rceil = B'' \tag{12}$$

$$i[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \Longrightarrow B''' \tag{13}$$

From (12), (10) and Lemma 3 (*Violation Detectability*) we obtain

$$i[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \xrightarrow{\text{fail!}} \tag{14}$$

and from (13), (14) and Lemma 4 (*Detection Preservation*) we get $B''' \xrightarrow{\text{fail!}}$. Hence, by (11), and standard transition rules for parallel composition and scoping, PAR and SCP, we can reconstruct $B' \xrightarrow{\text{fail!}}$, as required.

For the *if* case we assume:

$$(\nu i)(\lceil A \rceil \parallel i[\text{Mon}(\varphi)]^*) \xrightarrow{s} B \tag{15}$$

$$B \Downarrow_{\text{fail!}} \tag{16}$$

and have to prove $A, s \models_{\nu} \varphi$. From (16) we know $B \xrightarrow{\text{fail!}}$. Together with (15), this implies that there exists a sequence of reductions (τ -transitions) from $(\nu i)(\lceil A \rceil \parallel i[\text{Mon}(\varphi)]^*)$ leading to a (monitored) system that flags a violation:

$$\exists B' \text{ s.t. } (\nu i)(\lceil A \rceil \parallel i[\text{Mon}(\varphi)]^*) \xrightarrow{s} B' \xrightarrow{\text{fail!}} \tag{17}$$

From Lemma 5 (Monitor Separability) and (17) we obtain

$$\exists B'', B''' \text{ s.t. } B' \equiv (\nu i)(B'' \parallel B''') \tag{18}$$

$$A \xrightarrow{s} A' \text{ for some } A' \text{ where } \lceil A' \rceil = B'' \tag{19}$$

$$i[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \Longrightarrow B''' \tag{20}$$

From (17), (18) and the assumption that action fail! is fresh to A , we deduce that it can only be B''' that is capable of flagging the violation:

$$B''' \xrightarrow{\text{fail!}}. \tag{21}$$

Thus, by (20) and (21), we get

$$i[\text{Mon}(\varphi) \triangleleft \text{tr}(s)] \xrightarrow{\text{fail!}} \tag{22}$$

Therefore, by (19), (22) and Lemma 3 (Violation Detectability) we obtain $A, s \models_v \varphi$. \square

7 Proving the monitor sub-properties

The proof of monitor correctness given in Sect. 6 hinges on the three monitor sub-properties discussed in the same section. We now consider the proofs for these three monitor properties for the monitor synthesis presented in Sect. 5. In what follows, we assume that the tracer, i.e. where the synthesised monitor is placed in the instrumentation of Definition 6, is located at the process identifier i_{mtr} . Moreover sequences of identifiers are denoted as \tilde{h} ; for instance, $(\nu \tilde{j})A$ is used to denote the system $(\nu j_1) \dots (\nu j_n)A$ whenever the scoped identifiers are unimportant.

These proofs rely on an encoding of formula substitutions, $\theta : \text{LVAR} \rightarrow \text{SHML}$, partial maps from formula variables to (possibly open) formulas, to lists of tuples containing a string representation of the variable and the respective monitor translation of the formula as defined in Definition 7. Formula substitutions are denoted as lists of individual substitutions, $\{\varphi_1/X_1\} \dots \{\varphi_n/X_n\}$ where every X_i is distinct, and empty substitutions are denoted as ϵ .

Definition 8 (Formula Substitution Encoding)

$$\text{enc}(\theta) \stackrel{\text{def}}{=} \begin{cases} \text{nil} & \text{when } \theta = \epsilon \\ \{\langle X', \llbracket \varphi \rrbracket^{\text{m}} \rangle : \text{enc}(\theta')\} & \text{if } \theta = \{\max(X, \varphi)/X\}\theta' \end{cases}$$

Our monitor lookup function of Definition 7 models variable substitution, Lemma 6. We can also show that different representations of the same formula substitution do not affect the outcome of the execution of `lookUp` on the respective encoding, Lemma 7, which justifies the abuse of notation that assume a unique representation of a formula substitution.

Lemma 6 *If $\theta(X) = \varphi$ then $i[\text{lookUp}(X', \text{enc}(\theta)) \triangleleft q]^m \implies i[\llbracket \varphi \rrbracket^{\text{m}} \triangleleft q]^m$*

Proof By induction on the number of mappings $\{\varphi_1/X_1\} \dots \{\varphi_n/X_n\}$ in θ . The base case, i.e. when we have zero mappings, is trivial since it contradicts the left side of the implication. The inductive case has two subcases and follows from the pattern-matching branches of the `lookUp` code, defined in Definition 7. \square

Lemma 7 *If $\theta(X) = \varphi$ then $i[\text{lookUp}(X', \text{enc}(\theta')) \triangleleft q]^m \implies i[\llbracket \varphi \rrbracket^{\text{m}} \triangleleft q]^m$ whenever θ and θ' denote the same substitution.*

Proof By induction on the number of mappings $\{\varphi_1/X_1\} \dots \{\varphi_n/X_n\}$ in θ , analogously to the proof for Lemma 6. \square

Our proofs use another technical result, Lemma 8, stating that silent actions are, in some sense, preserved when actor-mailbox contents of a free actor are increased; note that the lemma only applies for cases where the mailbox at this free actor decreases in size or remains unaffected by the τ -action, specified through the sublist condition $q' \leq q$.

Lemma 8 (Mailbox Increase) $(\nu \tilde{h})(i[e \triangleleft q]^m \parallel A) \xrightarrow{\tau} (\nu \tilde{j})(i[e' \triangleleft q']^m \parallel B)$ where $i \notin \tilde{h}$ and $q' \leq q$ implies $(\nu \tilde{h})(i[e \triangleleft q : v]^m \parallel A) \xrightarrow{\tau} (\nu \tilde{j})(i[e' \triangleleft q' : v]^m \parallel B)$

Proof By rule induction on $(\nu \tilde{h})(i[e \triangleleft q]^m \parallel A) \xrightarrow{\tau} (\nu \tilde{j})(i[e' \triangleleft q']^m \parallel B)$. \square

7.1 Violation detection

In one direction, Lemma 3 (Violation Detection) from Sect. 6 relies on Lemma 13 (see Appendix A.1) in order to establish the correspondence between violations and the possibility of detections. In the other direction, Lemma 3 relies on Lemma 16 (see Appendix A.1), which establishes a correspondence between violation detections and actual violations, as stated in Definition 3. We recall that Lemma 3 was stated wrt. *closed* SHML formulas.

Lemma 3 (Violation Detection). *Whenever $A \xrightarrow{s}$ then :*

$$A, s \models_v \varphi \text{ iff } i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \xrightarrow{\text{fail!}}$$

Proof For the *only-if* case, we assume $A \xrightarrow{s}$ and $A, s \models_v \varphi$ and are required to prove $i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \xrightarrow{\text{fail!}}$. We recall from Sect. 5 that **Mon** was defined as

$$\lambda x_{\text{frm}}.z_{\text{pid}} = \text{spw} (\llbracket x_{\text{frm}} \rrbracket^{\text{m}}(\text{nil}), \text{mLoop}(z_{\text{pid}})). \tag{23}$$

and as a result we can deduce (using rules such as APP, SPW and PAR) that

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \implies (v i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \varphi \rrbracket^{\text{m}}(\text{nil})]^*) \tag{24}$$

Assumption $A, s \models_v \varphi$ can be rewritten as $A, s \models_v \varphi\theta$ for $\theta = \epsilon$, and thus, by Definition 8 we know $\text{nil} = \text{enc}(\theta)$. By Lemma 13 we obtain

$$(v i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \varphi \rrbracket^{\text{m}}(\text{nil})]^*) \xrightarrow{\text{fail!}} \tag{25}$$

and the result thus follows from (24) and (25).

For the *if* case, we assume $A \xrightarrow{s}$ and $i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \xrightarrow{\text{fail!}}$ and are required to prove $A, s \models_v \varphi$.

Since φ is closed, we can assume the empty list of substitutions $\theta = \epsilon$ where, by default, $\text{fv}(\varphi) \subseteq \text{dom}(\theta)$ and, by Definition 8, $\text{nil} = \text{enc}(\theta)$. By (23) we can decompose the transition sequence $i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \xrightarrow{\text{fail!}}$ as

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \xrightarrow{(\tau)}^3 (v i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \varphi \rrbracket^{\text{m}}(\text{nil})]^*) \xrightarrow{\text{fail!}} \tag{26}$$

The result, i.e. $A, s \models_v \varphi$, follows from (26) and Lemma 16. □

7.2 Detection preservation

In order to prove Lemma 4 from Sect. 6, we are able to require a stronger guarantee, i.e. confluence (Definition 9) under weak transitions for the concurrent monitors described in Definition 7; this is formalised as Lemma 18 in Appendix A.2.

Definition 9 (*Confluence modulo Inputs with Identical Recipients*)

$$\text{cnf}(A) \stackrel{\text{def}}{=} A \xrightarrow{\gamma_1} A' \text{ and } A \xrightarrow{\gamma_2} A'' \text{ implies } \begin{cases} \gamma_1 = i?v_1, \gamma_2 = i?v_2 \text{ or;} \\ \gamma_1 = \gamma_2, A' = A'' \text{ or;} \\ A' \xrightarrow{\gamma_2} A''', A'' \xrightarrow{\gamma_1} A''' \text{ for some } A''' \end{cases}$$

In Definition 9, a system is deemed *confluent* if, whenever it can perform two *separate* actions γ_1 and γ_2 that are *not* input actions at the same actor, both residual systems can *resp.* still perform the other action to reach the same system.

Lemma 18 (Appendix A.2) allows us to prove Lemma 9, and subsequently Lemma 10; the latter Lemma implies Detection Preservation, Lemma 4, used by Theorem 3 of Sect. 6.

Lemma 9 For all $\varphi \in \text{SHML}$, $q \in \text{VAL}^*$

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^* \Longrightarrow A, A \xrightarrow{\text{fail!}} \text{ and } A \xrightarrow{\tau} B \text{ implies } B \xrightarrow{\text{fail!}}$$

Proof From $i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^* \Longrightarrow A$ and Lemma 18 we know that $\text{cnf}(A)$. The proof is by induction on $A(\xrightarrow{\tau})^n \cdot \xrightarrow{\text{fail!}}$.

$n = 0$: We have $A \xrightarrow{\text{fail!}} A'$ (for some A'). By $A \xrightarrow{\tau} B$ and $\text{cnf}(A)$ we obtain $B \xrightarrow{\text{fail!}} B'$ for some B' where $A' \xrightarrow{\tau} B'$.

$n = k + 1$: We have $A \xrightarrow{\tau} A'(\xrightarrow{\tau})^k \cdot \xrightarrow{\text{fail!}}$ (for some A'). By $A \xrightarrow{\tau} A'$, $A \xrightarrow{\tau} B$ and $\text{cnf}(A)$ we either know that $B = A'$, in which case the result follows immediately, or else obtain

$$B \xrightarrow{\tau} A'' \tag{27}$$

$$A' \xrightarrow{\tau} A'' \text{ for some } A'' \tag{28}$$

In such a case, by $A \xrightarrow{\tau} A'$ and $i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^* \Longrightarrow A$ we deduce that

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^* \Longrightarrow A',$$

and subsequently, by (28), $A'(\xrightarrow{\tau})^k \cdot \xrightarrow{\text{fail!}}$ and I.H. we obtain $A'' \xrightarrow{\text{fail!}}$; the required result then follows from (27). \square

Lemma 10 (Detection Confluence) For all $\varphi \in \text{SHML}$, $q \in \text{VAL}^*$

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^* \Longrightarrow A, A \xrightarrow{\text{fail!}} \text{ and } A \Longrightarrow B \text{ implies } B \xrightarrow{\text{fail!}}$$

Proof By induction on $A(\xrightarrow{\tau})^n B$ and Lemma 9. \square

We are now in a position to prove Lemma 4 of Sect. 6.

Lemma 4 (Detection Preservation). For all $\varphi \in \text{SHML}$, $q \in \text{VAL}^*$

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^* \xrightarrow{\text{fail!}} \text{ and } i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^* \Longrightarrow B \text{ implies } B \xrightarrow{\text{fail!}}$$

Proof From Lemma 10, for the case where $i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^* \Longrightarrow i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^*$. \square

7.3 Monitor separability

For the proof for Lemma 5 of Sect. 6, we make use of Lemma 2, relating the behaviour of a monitored system to the same system when unmonitored, Lemma 8 delineating behaviour preservation after extending mailbox contents at specific actors, and Lemma 11 (Appendix A), so as to reason about the structure and generic behaviour of synthesised monitors.

Lemma 5 (Monitor Separability). For all basic actors $\varphi \in \text{SHML}$, $A \in \text{ACTR}$ where i_{mtr} is fresh to A , and $s \in (\text{ACT} \setminus \{\text{fail!}\})^*$,

$$(\nu i_{\text{mtr}})(\lceil A \rceil \parallel i_{\text{mtr}}[\text{Mon}(\varphi)]^*) \xrightarrow{s} B \text{ implies } \exists B', B'' \text{ s.t. } \begin{cases} B \equiv (\nu i_{\text{mtr}})(B' \parallel B'') \\ A \xrightarrow{s} A' \text{ s.t. } B' = \lceil A' \rceil \\ i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \Longrightarrow B'' \end{cases}$$

Proof By induction on n in $(\nu i_{\text{mtr}})(\lceil A \rceil \parallel i_{\text{mtr}}[\text{Mon}(\varphi)]^*)(\xrightarrow{\gamma^n})^n B$, the length of the sequence of actions:

$n = 0$: Since $s = \epsilon$ and $A = (\nu i_{\text{mtr}})(\lceil A \rceil \parallel i_{\text{mtr}}[\text{Mon}(\varphi)]^*)$, the conditions hold trivially.

$n = k + 1$: We have $(\nu i_{\text{mtr}})(\lceil A \rceil \parallel i_{\text{mtr}}[\text{Mon}(\varphi)]^*)(\xrightarrow{\gamma^k})^k C \xrightarrow{\gamma} B$. By I.H. we know that

$$C \equiv (\nu i_{\text{mtr}})(C' \parallel C'') \tag{29}$$

$$A \xrightarrow{t} A'' \text{ s.t. } C' = \lceil A'' \rceil \tag{30}$$

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(t)]^* \implies C'' \tag{31}$$

$$\gamma = \tau \text{ implies } t = s \text{ and } \gamma = \alpha \text{ implies } t\alpha = s \tag{32}$$

and by (31) and Lemma 11 we know that

$$C'' \equiv (\nu \mathbf{h})(i_{\text{mtr}}[e \triangleleft q]^* \parallel C''') \tag{33}$$

$$\text{fId}(C'') = \{i_{\text{mtr}}\} \tag{34}$$

We proceed by considering the two possible subcases for the structure of γ :

$\gamma = \alpha$: By (32) we know that $s = t\alpha$. By (34) and (33), it must be the case that $C \equiv (\nu i_{\text{mtr}})(C' \parallel C'') \xrightarrow{\alpha} B$ happens because

$$\text{for some } B' \ C' \xrightarrow{\alpha} B' \tag{35}$$

$$B \equiv (\nu i_{\text{mtr}})(B' \parallel (\nu \mathbf{h})(i_{\text{mtr}}[e \triangleleft q : \text{tr}(\alpha)]^* \parallel C''')) \tag{36}$$

By (35), (30) and Lemma 2 we know that $\exists A'$ such that $\lceil A' \rceil = B'$ and that $A'' \xrightarrow{\alpha} A'$. Thus by (30) and $s = t\alpha$ we obtain

$$A \xrightarrow{s} A' \text{ s.t. } B' = \lceil A' \rceil$$

By (31), (33) and repeated applications of Lemma 8 we also know that

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(t) : \text{tr}(\alpha)]^* = i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft \text{tr}(s)]^* \implies (\nu \mathbf{h})(i_{\text{mtr}}[e \triangleleft q : \text{tr}(\alpha)]^* \parallel C''') = B''$$

The result then follows from (36).

$\gamma = \tau$: Analogous to the other case, where we also have the case that the reduction is instigated by C'' , in which case the results follows immediately. \square

8 Conclusion

We have studied a more intensional notion of correctness for monitor synthesis in a concurrent online setting; we worked close to the actual implementation level of abstraction so as to enhance our confidence in the correctness of our instrumented monitors. More precisely, we have identified a number of additional issues raised when proving monitor correctness in this concurrent setting, and built a tool [21], automating monitor synthesis from a reactive property logic (SHML) to asynchronous monitors in a concurrent language (Erlang), illustrating these issues. The specific contributions of the paper, in order of importance, are:

1. A novel formal *definition of monitor correctness*, Definition 6, dealing with issues such as system non-determinism and system interference.
2. A *proof technique* teasing apart aspects of the monitor correctness definition, Lemma 3, Lemma 4 and Lemma 5, allowing us to prove correctness in stages. We subsequently apply this technique to prove the *correctness of our tool*, Theorem 3.
3. An alternative *violation* characterisation of the logic, SHML, that is more amenable to runtime analysis and reasoning about monitor correctness, together with a proof of correspondence for this reformulation, Theorem 1.
4. An extension of a formalisation for Erlang describing its tracing semantics, Sect. 2.
5. A formal monitor synthesis definition from SHML formulas to Erlang code, Definition 7.

8.1 Related work

The aforementioned work, [5,23,34], discusses monitor synthesis from a different logic, namely LTL, to either pseudocode, automata or Büchi automata; none of this work considers online concurrent monitoring, circumventing issues associated with concurrency and system interference. There is considerable work on runtime monitoring of web services, e.g. [7,18] verifying the correctness of reactive (communication) properties, similar to those expressed through SHML; to the best of our knowledge, none of this work tackles correct monitor synthesis from a specified logic. In [14], Colombo et al. develop an Erlang RV tool using the EVM tracing mechanism but do not consider the issue of correct monitor generation. Fredlund [22] adapted a variant of HML to specify correctness properties in Erlang, albeit for model checking purposes.

There is also work relating HML formulas with tests, namely [1], but also [10]. Our monitors differ from tests, as in [1], in a number of ways: (i) they are defined in terms of *concurrent* actors, as opposed to *sequential* CCS processes; (ii) they analyse systems *asynchronously*, acting on traces, whereas tests interact with the system *directly*, forcing certain system behaviour; (iii) they are expected to *always* detect violations when they occur whereas tests are only required to have *one* possible execution that detects violations.

In recent work, Bocchi et al. [6] studied monitors that enforce multi-party session types at runtime for high-level specifications of message-passing programs, expressed using a distributed π -calculus. Their methodology differs from ours in a number of ways. In particular, (i) they give a direct operational semantics to their session specifications in terms of an LTS, which allows them to interact directly with the processes that they monitor; by contrast, we synthesise monitors from our specification formulas as programs in the host language, and focus on proving the correctness of this synthesis; (ii) the parallelisation criteria for their session projections is based on the participants being monitored whereas we parallelise on the basis of the structure of the formula, namely across formula conjunctions; (iii) they work at the level of an abstract model, namely a distributed π -calculus, whereas we work at a level of abstraction that is close to actual Erlang code that can be compiled and executed.

There is other work on the decomposition of monitor synthesis. In [4], they synthesise a dedicated monitor from an LTL specification for each *synchronous* component executing in a system, thereby *localising* monitoring to a component level. By contrast, our monitor synthesis is agnostic to the internal structure of the monitored system, and decomposition is purely based on the structure of the correctness formulas. In [35], they define a distributed logic for specifying correctness properties of distributed systems and provide a distributed monitor synthesis algorithm for the logic, implemented as actor-based tool called DiAna. Their setting is however different from ours: their systems do not assume a global clock and monitoring work over partially ordered traces (one for each location). Monitoring in the

absence of global clocks is also considered in [20], where they develop bisimulation-based coinductive techniques to reason about monitored systems. Crucially, none of these works considers issues relating to the correctness of monitor synthesis studied in this paper.

8.2 Future work

The monitoring semantics of Sect. 2 can be used as a basis to prove the correctness of existing Erlang monitoring tools such as [14, 15]. SHML can also be extended to handle limited, monitorable forms of liveness properties (often termed co-safety properties [28]); the work carried out in [10] provides an ideal starting point. It is also worth exploring mechanisms for synchronous monitoring, as opposed to asynchronous variant studied in this paper. Erlang also facilitates monitor *distribution* which can be used to lower monitoring overheads [16, 35]. Distributed monitoring can also be used to increase the expressivity of our tool so as to handle correctness properties for distributed programs. However, this poses a departure from our setting because the unique trace described by our framework would be replaced by separate independent traces at each location, where the lack of a total ordering of events may prohibit the detection of certain violations [20].

Acknowledgments The research work disclosed in this publication is partially funded by the Strategic Educational Pathways Scholarship Scheme (Malta). The scholarship is part nanced by the European Union European Social Fund.

A Auxiliary proofs

For the proofs in Sect. 7, we find it convenient to prove a technical result, Lemma 11, identifying the possible structures a monitor can be in after an arbitrary number of silent actions; the lemma also establishes that the only possible external action that a synthesised monitors can perform is the *fail* action: this property helps us to reason about the possible interactions that concurrent monitors may engage in.

Lemma 11 (Monitor Transitions and Structure) *For all $\varphi \in \text{SHML}$, $q \in (\text{VAL})^*$, $\theta : : \text{LVAR} \rightarrow \text{SHML}$, if $i[\llbracket \varphi \rrbracket^{\text{m}}(\text{enc}(\theta)) \triangleleft q]^{\bullet} (\rightarrow)^n A$ then*

1. $A \xrightarrow{\alpha} B$ implies $\alpha = \text{fail!}$ and;
2. A has the form $i[\llbracket \varphi \rrbracket^{\text{m}}(\text{enc}(\theta)) \triangleleft q]^{\bullet}$ or, depending on φ :

$$\begin{aligned} \varphi = \text{ff}: & A \equiv i[\text{fail!} \triangleleft q]^{\bullet} \text{ or } A \equiv i[\text{fail} \triangleleft q]^{\bullet} \\ \varphi = [\alpha]\psi: & A \equiv i[\text{rcv}(\text{tr}(\alpha) \rightarrow \llbracket \psi \rrbracket^{\text{m}}(\text{enc}(\theta)); _ \rightarrow \text{ok}) \text{end} \triangleleft q]^{\bullet} \text{ or} \\ & (A \equiv B \text{ where } i[\llbracket \psi \rrbracket^{\text{m}}(\text{enc}(\theta)) \triangleleft r]^{\bullet} (\xrightarrow{\tau})^k B \text{ for some } k < n \text{ and } q = \text{tr}(\alpha) : r) \\ & \text{or} \\ & A \equiv i[\text{ok} \triangleleft r]^{\bullet} \text{ where } q = u : r \\ \varphi = \varphi_1 \wedge \varphi_2: & A \equiv i \left[\begin{array}{l} y_1 = \text{spw}(\llbracket \varphi_1 \rrbracket^{\text{m}}(\text{enc}(\theta))), \\ y_2 = \text{spw}(\llbracket \varphi_2 \rrbracket^{\text{m}}(\text{enc}(\theta))), \text{fork}(y_1, y_2) \triangleleft q \end{array} \right]^{\bullet} \\ & \text{or} \\ & A \equiv (v j_1) (i[e \triangleleft q]^{\bullet} \parallel (v \tilde{h}_1)(j_1[e_1 \triangleleft q_1]^{\bullet} \parallel B)) \text{ where} \\ & \quad - e \text{ is } y_1 = j_1, y_2 = \text{spw}(\llbracket \varphi_2 \rrbracket^{\text{m}}(\text{enc}(\theta))), \text{fork}(y_1, y_2) \text{ or} \\ & \quad y_2 = \text{spw}(\llbracket \varphi_2 \rrbracket^{\text{m}}(\text{enc}(\theta))), \text{fork}(j_1, y_2) \\ & \quad - j_1[\llbracket \varphi_1 \rrbracket^{\text{m}}(\text{enc}(\theta))]^{\bullet} (\xrightarrow{\tau})^k (v \tilde{h}_1)(j_1[e_1 \triangleleft q_1]^{\bullet} \parallel B) \text{ for some } k < n \\ & \text{or} \end{aligned}$$

$$A \equiv (\nu j_1, j_2) \left(i[y_2 = j_2, \mathit{fork}(j_1, y_2) \triangleleft q]^\bullet \parallel (\nu \tilde{h}_1)(j_1[e_1 \triangleleft q_1]^\bullet \parallel B) \parallel (\nu \tilde{h}_2)(j_2[e_2 \triangleleft q_2]^\bullet \parallel C) \right)$$

where

$$\begin{aligned} & - j_1[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta))]^\bullet (\xrightarrow{\tau})^k (\nu \tilde{h}_1)(j_1[e_1 \triangleleft q_1]^\bullet \parallel B) \text{ for some } k < n \\ & - j_2[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta))]^\bullet (\xrightarrow{\tau})^l (\nu \tilde{h}_2)(j_2[e_2 \triangleleft q_2]^\bullet \parallel C) \text{ for some } l < n \end{aligned}$$

or

$$A \equiv (\nu j_1, j_2) (i[e \triangleleft r]^\bullet \parallel (\nu \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B) \parallel (\nu \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C))$$

where

$$\begin{aligned} & - e \text{ is either } \mathit{fork}(j_1, j_2) \text{ or } (\mathit{rcv } z \rightarrow j_1!z, j_2!z \mathit{end}, \mathit{fork}(j_1, j_2)) \\ & \text{ or } j_1!u, i_2!u, \mathit{fork}(j_1, j_2) \text{ or } j_2!u, \mathit{fork}(j_1, j_2) \\ & - j_1[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta)) \triangleleft q_1]^\bullet (\xrightarrow{\tau})^k (\nu \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B) \text{ for } k < n, q_1 < q \\ & - j_2[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta)) \triangleleft q_2]^\bullet (\xrightarrow{\tau})^l (\nu \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C) \text{ for } l < n, q_2 < q \end{aligned}$$

$\varphi = X: A \equiv i[y = \mathit{lookUp}(X', \mathit{enc}(\theta')), y(\mathit{enc}(\theta)) \triangleleft q]^\bullet$ where $\theta' < \theta$ or

$$A \equiv i \left[y = \left(\begin{array}{l} \mathit{case } \mathit{enc}(\theta') \text{ of } \{X', z_{\mathit{mon}}\} : _ \rightarrow z_{\mathit{mon}}; \\ _ : z_{i_l} \rightarrow \mathit{lookUp}(X', z_{i_l}); \\ \mathit{nil} \rightarrow \mathit{exit}; \\ \mathit{end} \end{array} \right), y(\mathit{enc}(\theta)) \triangleleft q \right]^\bullet$$

where $\theta' < \theta$, or

$A \equiv B$ where

$$\begin{aligned} & - i[y = \llbracket \psi \rrbracket^{\mathbf{m}}, y(\mathit{enc}(\theta)) \triangleleft q]^\bullet (\xrightarrow{\tau})^k B \\ & - \theta(X) = \psi \text{ or } A \equiv i[y = \mathit{exit}, y(\mathit{enc}(\theta)) \triangleleft q]^\bullet \text{ or } A \equiv i[\mathit{exit} \triangleleft q]^\bullet \end{aligned}$$

$\varphi = \max(X, \psi): A \equiv B$ where $i[\llbracket \psi \rrbracket^{\mathbf{m}}(X', \llbracket \psi \rrbracket^{\mathbf{m}}) : \mathit{enc}(\theta) \triangleleft q]^\bullet (\xrightarrow{\tau})^k B$ for $k < n$.

Proof The proof is by strong induction on $i[\llbracket \varphi \rrbracket^{\mathbf{m}}(I_{\mathit{env}}) \triangleleft q]^\bullet (\xrightarrow{\tau})^n A$. The inductive case involves a long and tedious list of case analysis exhausting all possibilities. \square

A.1 Proofs for establishing violation detection

Lemma 13 uses Lemma 12 which relates possible detections by monitors synthesised from subformulas to possible detections by monitors synthesised from conjunctions using these subformulas.

Lemma 12 For an arbitrary θ , $(\nu i)(i_{\mathit{mtr}}[\mathit{mLoop}(j_1) \triangleleft \mathit{tr}(s)]^* \parallel i[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta))]^\bullet) \xrightarrow{\mathit{fail}!}$ implies $(\nu i)(i_{\mathit{mtr}}[\mathit{mLoop}(i) \triangleleft \mathit{tr}(s)]^* \parallel i[\llbracket \varphi_1 \wedge \varphi_2 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta))]^\bullet) \xrightarrow{\mathit{fail}!}$ for any $\varphi_2 \in \mathit{SHML}$.

Proof By Definition 7, we know that we can derive the sequence of reductions

$$\begin{aligned} & (\nu i)(i_{\mathit{mtr}}[\mathit{mLoop}(i) \triangleleft \mathit{tr}(s)]^* \parallel i[\llbracket \varphi_1 \wedge \varphi_2 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta))]^\bullet) \implies \\ & (\nu i)(i_{\mathit{mtr}}[\mathit{mLoop}(i) \triangleleft \mathit{tr}(s)]^* \parallel (\nu j, h)(i[\mathit{fork}(j, h)]^\bullet \parallel j[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta))]^\bullet \parallel h[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta))]^\bullet)) \end{aligned}$$

We then prove, by induction on the structure of s , the following (see [19] for details):

$$\begin{aligned} & (\nu i)(i_{\mathit{mtr}}[\mathit{mLoop}(i) \triangleleft \mathit{tr}(s)]^\bullet \parallel i[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta)) \triangleleft q]^\bullet) \xrightarrow{\mathit{fail}!} \text{ implies} \\ & (\nu i) \left(i_{\mathit{mtr}}[\mathit{mLoop}(i) \triangleleft \mathit{tr}(s)]^* \parallel (\nu j, h)(i[\mathit{fork}(j, h)]^\bullet \parallel j[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta)) \triangleleft q]^\bullet \parallel h[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(\mathit{enc}(\theta)) \triangleleft q]^\bullet) \right) \xrightarrow{\mathit{fail}!} \end{aligned}$$

\square

Lemma 13 *If $A, s \models_v \varphi\theta$ and $l_{env} = enc(\theta)$ then*

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(s)]^* \parallel i[[\varphi]^m(l_{env})]^\bullet) \xrightarrow{fail!}.$$

Proof Proof by rule induction on $A, s \models_v \varphi\theta$:

$A, s \models_v \mathbf{ff}\theta$: Using Definition 7 for the definition of $[[\mathbf{ff}]^m]$ and the rule APP (and PAR and SCP), we have

$$\begin{aligned} &(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(s)]^* \parallel i[[\mathbf{ff}]^m(l_{env})]^\bullet) \\ &\implies (\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(s)]^* \parallel i[\mathbf{fail!}]^\bullet) \end{aligned}$$

The result follows trivially, since the process i can transition with a $\mathbf{fail!}$ action in a single step using the rule SNDU.

$A, s \models_v (\varphi_1 \wedge \varphi_2)\theta$ because $A, s \models_v \varphi_1\theta$: By $A, s \models_v \varphi_1\theta$ and I.H. we have

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(s)]^* \parallel i[[\varphi_1]^m(l_{env})]^\bullet) \xrightarrow{fail!}$$

The result thus follows from Lemma 12, which allows us to conclude that

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(s)]^* \parallel i[[\varphi_1 \wedge \varphi_2]^m(l_{env})]^\bullet) \xrightarrow{fail!}$$

$A, s \models_v (\varphi_1 \wedge \varphi_2)\theta$ because $A, s \models_v \varphi_2\theta$: Analogous.

$A, s \models_v ([\alpha]\varphi)\theta$ because $s = \alpha t, A \xrightarrow{\alpha} B$ and $B, t \models_v \varphi\theta$: Using the rule APP SCP and Definition 7 for the property $[\alpha]\varphi$ we derive (37), by executing \mathbf{mLoop} — see Definition 7 — we obtain (38), and then by rule RD1 we derive (39) below.

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(\alpha t)]^* \parallel i[[\varphi]^m(l_{env})]^\bullet) \xrightarrow{\tau} \tag{37}$$

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(\alpha t)]^* \parallel i[\mathbf{rcv}(tr(\alpha) \rightarrow [\varphi]^m(l_{env}); _ \rightarrow \mathbf{ok})\mathbf{end}]^\bullet) \implies \tag{38}$$

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(t)]^* \parallel i[\mathbf{rcv}(tr(\alpha) \rightarrow [\varphi]^m(l_{env}); _ \rightarrow \mathbf{ok})\mathbf{end} \triangleleft tr(\alpha)]^\bullet) \xrightarrow{\tau} \tag{39}$$

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(t)]^* \parallel i[[\varphi]^m(l_{env})]^\bullet)$$

By $B, t \models_v \varphi\theta$ and I.H. we obtain

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(t)]^* \parallel i[[\varphi]^m(l_{env})]^\bullet) \xrightarrow{fail!}$$

and, thus, the result follows by (37), (38) and (39).

$A, s \models_v (\mathbf{max}(X, \varphi))\theta$ because $A, s \models_v \varphi\{\mathbf{max}(X, \varphi)/X\}\theta$: By Definition 7 and APP for process i , we derive

$$\begin{aligned} &(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(s)]^* \parallel i[[\mathbf{max}(X, \varphi)]^m(l_{env})]^\bullet) \implies \\ &(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(s)]^* \parallel i[[\varphi]^m(\{X', [\varphi]^m\}: l_{env})]^\bullet) \end{aligned} \tag{40}$$

Assuming the appropriate α -conversion for X in $\mathbf{max}(X, \varphi)$, we note that from $l_{env} = enc(\theta)$ and Definition 8 we obtain

$$enc(\{\mathbf{max}(X, \varphi)/X\}\theta) = \{X', [\varphi]^m\}: l_{env} \tag{41}$$

By $A, s \models_v \varphi\{\mathbf{max}(X, \varphi)/X\}\rho$, (41) and I.H. we obtain

$$(\nu i)(i_{mtr}[\mathbf{mLoop}(i) \triangleleft tr(s)]^* \parallel i[[\varphi]^m(\{X', [\varphi]^m\}: l_{env})]^\bullet) \xrightarrow{fail!} \tag{42}$$

The result follows from (40) and (42). □

Lemma 16 relies on a technical result, Lemma 15 which allows us to recover a violating reduction sequence for a subformula φ_1 or φ_2 from that of the synthesised monitor of a conjunction formula $\varphi_1 \wedge \varphi_2$. Lemma 15 relies on Lemma 14.

Lemma 14 *For some $l \leq n$:*

$$\begin{aligned}
 & (\nu j, h) \left(i \left[\text{fork}(j, h) \triangleleft q_{\text{frk}} \right]^{\bullet} \parallel j \llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}}) \triangleleft q \right]^{\bullet} \parallel h \llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}}) \triangleleft r \right]^{\bullet} \xrightarrow{(\tau)}^n \text{fail!} \\
 & \text{implies } (\nu j) (i_{\text{mtr}}[\text{mLoop}(j) \triangleleft q_{\text{frk}}]^* \parallel j \llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}}) \triangleleft q \right]^{\bullet} \xrightarrow{(\tau)}^l \text{fail!} \\
 & \text{or } (\nu h) (i_{\text{mtr}}[\text{mLoop}(h) \triangleleft q_{\text{frk}}]^* \parallel h \llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}}) \triangleleft r \right]^{\bullet} \xrightarrow{(\tau)}^l \text{fail!}
 \end{aligned}$$

Proof By induction on the structure of the mailbox q_{frk} at actor i . □

Lemma 15 *For some $l \leq n$*

$$\begin{aligned}
 & (\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel (\nu j, h) \left(i \left[\text{fork}(j, h) \triangleleft \text{tr}(t) \right]^{\bullet} \parallel j \llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \parallel h \llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \right) \right) \\
 & \xrightarrow{(\tau)}^k \text{fail!} \\
 & \text{implies } (\nu i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(ts)]^* \parallel i \llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \xrightarrow{(\tau)}^l \text{fail!} \\
 & \text{or } (\nu i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(ts)]^* \parallel i \llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \xrightarrow{(\tau)}^l \text{fail!}
 \end{aligned}$$

Proof Proof by induction on the structure of s .

$s = \epsilon$: From the structure of mLoop , we know that after the function application, the actor $i_{\text{mtr}}[\text{mLoop}(i)]^*$ is stuck. Thus we conclude that it must be the case that

$$(\nu j, h) \left(i \left[\text{fork}(j, h) \triangleleft \text{tr}(t) \right]^{\bullet} \parallel j \llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \parallel h \llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \right) \xrightarrow{(\tau)}^k \text{fail!}$$

where $k = n$ or $k = n - 1$. In either case, the required result follows from Lemma 14.
 $s = \alpha s'$: We have two subcases:

If

$$(\nu j, h) \left(i \left[\text{fork}(j, h) \triangleleft \text{tr}(t) \right]^{\bullet} \parallel j \llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \parallel h \llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \right) \xrightarrow{(\tau)}^k \text{fail!}$$

for some $k \leq n$ then, by Lemma 14 we obtain

$$\begin{aligned}
 & (\nu j) (i_{\text{mtr}}[\text{mLoop}(j) \triangleleft \text{tr}(t)]^* \parallel j \llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \xrightarrow{(\tau)}^l \text{fail!} \\
 & \text{or } (\nu h) (i_{\text{mtr}}[\text{mLoop}(h) \triangleleft \text{tr}(t)]^* \parallel h \llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \xrightarrow{(\tau)}^l \text{fail!}
 \end{aligned}$$

for some $l \leq k$. By Lemma 8 we thus obtain

$$\begin{aligned}
 & (\nu j) (i_{\text{mtr}}[\text{mLoop}(j) \triangleleft \text{tr}(ts)]^* \parallel j \llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \xrightarrow{(\tau)}^l \text{fail!} \\
 & \text{or } (\nu h) (i_{\text{mtr}}[\text{mLoop}(h) \triangleleft \text{tr}(ts)]^* \parallel h \llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \xrightarrow{(\tau)}^l \text{fail!}
 \end{aligned}$$

as required.

Otherwise, it must be the case that

$$(\nu i) \left(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel (\nu j, h) \left(i [\text{fork}(j, h) \triangleleft \text{tr}(t)]^\bullet \parallel j[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet \parallel h[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet \right) \right) (\xrightarrow{\tau})^k \quad (43)$$

$$(\nu i) \left(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s')]^* \parallel (\nu j, h) (i [e_{\text{fork}} \triangleleft q : \text{tr}(\alpha)]^\bullet \parallel A) \right) (\xrightarrow{\tau})^{n-k} \xrightarrow{\text{fail}!} \quad (44)$$

For some $k = 3 + k_1$ where

$$(\nu j, h) \left(i [\text{fork}(j, h) \triangleleft \text{tr}(t)]^\bullet \parallel j[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet \parallel h[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet \right) (\xrightarrow{\tau})^{k_1} \quad (45)$$

$$(\nu j, h) (i [e_{\text{fork}} \triangleleft q]^\bullet \parallel A)$$

By (45) and Lemma 8 we obtain

$$(\nu j, h) \left(i [\text{fork}(j, h) \triangleleft \text{tr}(t) : \text{tr}(\alpha)]^\bullet \parallel j[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet \parallel h[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet \right) (\xrightarrow{\tau})^{k_1}$$

$$(\nu j, h) (i [e_{\text{fork}} \triangleleft q : \text{tr}(\alpha)]^\bullet \parallel A)$$

and by (44) we can construct the sequence of transitions:

$$(\nu i) \left(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s')]^* \parallel (\nu j, h) \left(i [\text{fork}(j, h) \triangleleft \text{tr}(t) : \alpha]^\bullet \parallel j[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet \parallel h[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet \right) \right) (\xrightarrow{\tau})^{n-3} \xrightarrow{\text{fail}!}$$

Thus, by I.H. we obtain, for some $l \leq n - 3$

$$(\nu i) (i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(t\alpha s')]^* \parallel i[\llbracket \varphi_1 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) (\xrightarrow{\tau})^l \xrightarrow{\text{fail}!}$$

$$\text{or } (\nu i) (i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(t\alpha s')]^* \parallel i[\llbracket \varphi_2 \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) (\xrightarrow{\tau})^l \xrightarrow{\text{fail}!}$$

The result follows since $s = \alpha s'$. □

Equipped with Lemma 15, we can now prove Lemma 16.

Lemma 16 *If $A \xrightarrow{s} \text{enc}(\theta)$ and $(\nu i) (i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \varphi \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) \xrightarrow{\text{fail}!}$ then $A, s \models_{\nu} \varphi\theta$, whenever $\text{fv}(\varphi) \subseteq \text{dom}(\theta)$.*

Proof By strong induction on the number of transitions n , leading to the action fail!

$$(\nu i) (i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \varphi \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) (\xrightarrow{\tau})^n \xrightarrow{\text{fail}!}$$

$n = 0$: By inspection of the definition for \mathbf{mLoop} , and by case analysis of $\llbracket \varphi \rrbracket^{\mathbf{m}}(l_{\text{env}})$ from Definition 7, it can never be the case that

$$(\nu i) (i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \varphi \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) \xrightarrow{\text{fail}!}.$$

Thus the result holds trivially.

$n = k + 1$: We proceed by case analysis on φ .

$\varphi = \text{ff}$: The result holds immediately for any A and s by Definition 3.

$\varphi = [\alpha]\psi$: By Definition 7, we know that

$$(\nu i)(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[[[\alpha]\psi]^{\text{m}}(l_{\text{env}})]^\bullet)(\xrightarrow{\tau})^{k_1} \tag{46}$$

$$(\nu i)(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_2)]^* \parallel i[[[\alpha]\psi]^{\text{m}}(l_{\text{env}}) \triangleleft \text{tr}(s_1)]^\bullet) \xrightarrow{\tau} \tag{47}$$

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_2)]^* \parallel i \left[\text{rcv} \left(\begin{array}{l} \text{tr}(\alpha) \rightarrow [[\psi]^{\text{m}}(l_{\text{env}})]^\bullet \\ _ \rightarrow \text{ok} \end{array} \right); \text{end} \triangleleft \text{tr}(s_1) \right]^\bullet \right) (\xrightarrow{\tau})^{k_2} \xrightarrow{\text{fail!}} \tag{48}$$

$$\text{where } k + 1 = k_1 + k_2 + 1 \text{ and } s = s_1 s_2 \tag{49}$$

From the analysis of the code in (48), the only way for the action fail! to be triggered is by choosing the guarded branch $\text{tr}(\alpha) \rightarrow [[\varphi]^{\text{m}}(l_{\text{env}})]^\bullet$ in actor i . This means that (48) can be decomposed into the following reduction sequences.

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_2)]^* \parallel i \left[\text{rcv} \left(\text{tr}(\alpha) \rightarrow [[\psi]^{\text{m}}(l_{\text{env}})]^\bullet; _ \rightarrow \text{ok} \right) \text{end} \triangleleft \text{tr}(s_1) \right]^\bullet \right) (\xrightarrow{\tau})^{k_3} \tag{50}$$

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_4)]^* \parallel i \left[\text{rcv} \left(\text{tr}(\alpha) \rightarrow [[\psi]^{\text{m}}(l_{\text{env}})]^\bullet; _ \rightarrow \text{ok} \right) \text{end} \triangleleft \text{tr}(s_1 s_3) \right]^\bullet \right) \xrightarrow{\tau} \tag{51}$$

$$(\nu i) i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_4)]^* \parallel i \left[[[\psi]^{\text{m}}(l_{\text{env}}) \triangleleft \text{tr}(s_5)]^\bullet \right] (\xrightarrow{\tau})^{k_4} \xrightarrow{\text{fail!}} \tag{52}$$

$$\text{where } k_2 = k_3 + k_4 + 1 \text{ and } s_1 s_3 = \alpha s_5 \text{ and } s_2 = s_3 s_4 \tag{53}$$

By (49) and (53) we derive

$$s = \alpha t \text{ where } t = s_5 s_4 \tag{54}$$

From the definition of mLoop we can derive

$$\begin{aligned} & (\nu i)(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(t)]^* \parallel i[[[\psi]^{\text{m}}(l_{\text{env}})]^\bullet])(\xrightarrow{\tau})^{k_5} \\ & (\nu i)(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_4)]^* \parallel i \left[[[\psi]^{\text{m}}(l_{\text{env}}) \triangleleft \text{tr}(s_5)]^\bullet \right]) \end{aligned} \tag{55}$$

where $k_5 \leq k_1 + k_3$. From (54) we can split $A \xrightarrow{s}$ as $A \xrightarrow{\alpha} A' \xrightarrow{t}$ and from (55), (52), the fact that $k_5 + k_4 < k + 1 = n$ from (49) and (53), and I.H. we obtain

$$A', t \models_v \psi \theta \tag{56}$$

From (56), $A \xrightarrow{\alpha} A'$ and Definition 3 we thus conclude $A, s \models_v ([\alpha]\psi)\theta$.

$\varphi = \varphi_1 \wedge \varphi_2$ From Definition 7, we can decompose the transition sequence as follows

$$(\nu i)(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[[[\varphi_1 \wedge \varphi_2]^{\text{m}}(l_{\text{env}})]^\bullet])(\xrightarrow{\tau})^{k_1} \tag{57}$$

$$(\nu i)(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_2)]^* \parallel i[[[\varphi_1 \wedge \varphi_2]^{\text{m}}(l_{\text{env}}) \triangleleft \text{tr}(s_1)]^\bullet]) \xrightarrow{\tau} \tag{58}$$

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_2)]^* \parallel i \left[\begin{array}{l} y_1 = \text{spw} \left([[\varphi_1]^{\text{m}}(l_{\text{env}})]^\bullet \right), \\ y_2 = \text{spw} \left([[\varphi_2]^{\text{m}}(l_{\text{env}})]^\bullet \right), \text{fork}(y_1, y_2) \triangleleft \text{tr}(s_1) \end{array} \right]^\bullet \right) (\xrightarrow{\tau})^{k_2} \tag{59}$$

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_4)]^* \parallel i \left[\begin{array}{l} y_1 = \text{spw} \left([[\varphi_1]^{\text{m}}(l_{\text{env}})]^\bullet \right), \\ y_2 = \text{spw} \left([[\varphi_2]^{\text{m}}(l_{\text{env}})]^\bullet \right), \text{fork}(y_1, y_2) \triangleleft \text{tr}(s_1 s_3) \end{array} \right]^\bullet \right) (\xrightarrow{\tau})^2 \tag{60}$$

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_4)]^* \parallel (\nu j) \left(i \left[\begin{array}{l} y_2 = \text{spw} (\llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})), \\ \text{fork}(j, y_2) \end{array} \right] \triangleleft \text{tr}(s_1 s_3) \right)^{\bullet} \right) \right) (\xrightarrow{\tau})^{k_3} \text{fail!} \tag{61}$$

where $k + 1 = k_1 + 1 + k_2 + 2 + k_3$, $s = s_1 s_2$ and $s_2 = s_3 s_4$ (62)

From (61) we can deduce that there are two possible transition sequences how action fail! was reached:

1. If fail! was reached because $j[\llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} (\xrightarrow{\tau})^{k_4} \text{fail!}$ on its own, for some $k_4 \leq k_3$ then, by PAR and SCP we deduce

$$(\nu i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel j[\llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet}) (\xrightarrow{\tau})^{k_4} \text{fail!}$$

From (62) we know that $k_4 < k + 1 = n$, and by the premise $A \xrightarrow{s}$ and I.H. we obtain $A, s \models_{\nu} \varphi_1 \theta$. By Definition 3 we then obtain $A, s \models_{\nu} (\varphi_1 \wedge \varphi_2) \theta$

2. Alternatively, (61) can be decomposed further as

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_4)]^* \parallel (\nu j) \left(i \left[\begin{array}{l} y_2 = \text{spw} (\llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})), \\ \text{fork}(j, y_2) \end{array} \right] \triangleleft \text{tr}(s_1 s_3) \right)^{\bullet} \right) \right) (\xrightarrow{\tau})^{k_4} \tag{63}$$

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_6)]^* \parallel (\nu j) \left(i \left[\begin{array}{l} y_2 = \text{spw} (\llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})), \\ \text{fork}(j, y_2) \end{array} \right] \triangleleft \text{tr}(s_1 s_3 s_5) \right)^{\bullet} \right) \right) (\xrightarrow{\tau})^2 \tag{64}$$

$$(\nu i) \left(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_6)]^* \parallel (\nu j, h) \left(i \left[\text{fork}(j, h) \triangleleft \text{tr}(s_1 s_3 s_5) \right]^{\bullet} \parallel j[\llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \parallel h[\llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet} \right) \right) \right) (\xrightarrow{\tau})^{k_5} \text{fail!} \tag{65}$$

where $k_3 = k_4 + 2 + k_5$ and $s_4 = s_5 s_6$ (66)

From (65) and Lemma 15 we know that, for some $k_6 \leq k_5$ either

$$(\nu i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_1 s_3 s_5 s_6)]^* \parallel i[\llbracket \varphi_1 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet}) (\xrightarrow{\tau})^{k_6} \text{fail!}$$

or $(\nu i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s_1 s_3 s_5 s_6)]^* \parallel i[\llbracket \varphi_2 \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet}) (\xrightarrow{\tau})^{k_6} \text{fail!}$

From (62) and (66) we know that $s = s_1 s_3 s_5 s_6$ and that $k_6 < k + 1 = n$. By I.H., we obtain either $A, s \models_{\nu} \varphi_1 \theta$ or $A, s \models_{\nu} \varphi_2 \theta$ and, in either case, by Definition 3 we deduce $A, s \models_{\nu} (\varphi_1 \wedge \varphi_2) \theta$.

$\varphi = X$ By Definition 7, we can deconstruct

$$(\nu i) (i_{\text{mtr}}[\text{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket X \rrbracket^{\text{m}}(l_{\text{env}})]^{\bullet}) (\xrightarrow{\tau})^{k+1} \text{fail!}$$

as

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket X \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) \Longrightarrow \xrightarrow{\tau} \tag{67}$$

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s_2)]^* \parallel i[y = \text{lookUp}(X', l_{\text{env}}), y(l_{\text{env}}) \triangleleft \text{tr}(s_1)]^\bullet) \Longrightarrow \xrightarrow{\tau} \tag{68}$$

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s_4)]^* \parallel i[y = v, y(l_{\text{env}}) \triangleleft \text{tr}(s_1s_3)]^\bullet) \Longrightarrow \xrightarrow{\tau} \tag{69}$$

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s_6)]^* \parallel i[v(l_{\text{env}}) \triangleleft \text{tr}(s_1s_3s_5)]^\bullet) \Longrightarrow \xrightarrow{\text{fail!}} \tag{70}$$

where $s = s_1s_2$, $s_2 = s_3s_4$ and $s_4 = s_5s_6$

Since $X \in \text{dom}(\theta)$, we know that $\theta(X) = \psi$ for some ψ . By the assumption $l_{\text{env}} = \text{enc}(\theta)$ and Lemma 6 we obtain that $v = \llbracket \psi \rrbracket^{\mathbf{m}}$. Hence, by (67), (68), (69) and (70) we can reconstruct

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \psi \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) (\xrightarrow{\tau})^{k_1} (v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s_6)]^* \parallel i[\llbracket \psi \rrbracket^{\mathbf{m}}(l_{\text{env}}) \triangleleft \text{tr}(s_1s_3s_5)]^\bullet) (\xrightarrow{\tau})^{k_2} \xrightarrow{\text{fail!}} \tag{71}$$

where $k_1 + k_2 < k + 1 = n$. By (71) and I.H. we obtain $A, s \models_v \psi$, which is the result required, since by $\theta(X) = \psi$ we know that $X\theta = \psi$.

$\varphi = \max(X, \psi)$ By Definition 7, we can deconstruct

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \max(X, \psi) \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) (\xrightarrow{\tau})^{k+1} \xrightarrow{\text{fail!}}$$

as follows:

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \max(X, \psi) \rrbracket^{\mathbf{m}}(l_{\text{env}})]^\bullet) (\xrightarrow{\tau})^{k_1} \xrightarrow{\tau}$$

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s_2)]^* \parallel i[\llbracket \psi \rrbracket^{\mathbf{m}}(\{X', \psi\}: l_{\text{env}}) \triangleleft \text{tr}(s_1)]^\bullet) (\xrightarrow{\tau})^{k_2} \xrightarrow{\text{fail!}}$$

from which we can reconstruct the transition sequence

$$(v\ i)(i_{\text{mtr}}[\mathbf{mLoop}(i) \triangleleft \text{tr}(s)]^* \parallel i[\llbracket \psi \rrbracket^{\mathbf{m}}(\{X', \psi\}: l_{\text{env}})]^\bullet) (\xrightarrow{\tau})^{k_1+k_2} \xrightarrow{\text{fail!}} \tag{72}$$

By the assumption $l_{\text{env}} = \Gamma(\theta)$ we deduce that $\{X', \psi\}: l_{\text{env}} = \text{enc}(\{\max(X, \psi)\}\theta)$ and, since $k_1 + k_2 < k + 1 = n$, we can use (72), $A \xrightarrow{s}$ and I.H. to obtain $A, s \models_v \psi\{\max(X, \psi)\}\theta$. By Definition 3 we then conclude $A, s \models_v \max(X, \psi)\theta$. \square

A.2 Proofs for establishing Detection Preservation

Lemma 18 relies heavily on Lemma 17.

Lemma 17 (Translation Confluence) *For all $\varphi \in \text{SHML}$, $q \in (\text{VAL})^*$ and $\theta : \text{LVAR} \rightarrow \text{SHML}$, $i[\llbracket \varphi \rrbracket^{\mathbf{m}}(\text{enc}(\theta)) \triangleleft q]^\bullet \Longrightarrow A$ implies $\text{cnf}(A)$.*

Proof Proof by strong numerical induction on n in $i[\llbracket \varphi \rrbracket^{\mathbf{m}}(\text{enc}(\theta)) \triangleleft q]^\bullet (\xrightarrow{\tau})^n A$.

$n = 0$: The only possible τ -action that can be performed by $i[\llbracket \varphi \rrbracket^{\mathbf{m}}(\text{enc}(\theta)) \triangleleft q]^\bullet$ is that for the function application of the monitor definition, i.e.

$$i[\llbracket \varphi \rrbracket^{\mathbf{m}}(\text{enc}(\theta)) \triangleleft q]^\bullet \xrightarrow{\tau} i[e \triangleleft q]^\bullet \text{ for some } e. \tag{73}$$

Apart from that $i[[\varphi]]^m(\text{enc}(\theta)) \triangleleft q]^\bullet$ can also only perform input action at i , i.e.

$$i[[\varphi]]^m(\text{enc}(\theta)) \triangleleft q]^\bullet \xrightarrow{i?v} i[[\varphi]]^m(\text{enc}(\theta)) \triangleleft q : v]^\bullet$$

On the one hand, we can derive $i[e \triangleleft q]^\bullet \xrightarrow{i?v} i[e \triangleleft q : v]^\bullet$. Moreover, from (73) and Lemma 8 we can deduce $i[[\varphi]]^m(\text{enc}(\theta)) \triangleleft q : v]^\bullet \xrightarrow{\tau} i[e \triangleleft q : v]^\bullet$ which allows us to close the confluence diamond.

$n = k + 1$: We proceed by case analysis on the property φ , using Lemma 11 to infer the possible structures of the resulting process. Again, most involving cases are those for conjunction translations, as they generate more than one concurrent actor; we discuss one of these below:

$\varphi = \varphi_1 \wedge \varphi_2$: By Lemma 11, A can have any of 4 general structures, one of which is

$$A \equiv (v \ j_1, j_2) \left(i[j_2!u, \text{fork}(j_1, j_2) \triangleleft q]^\bullet \parallel (v \ \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B) \parallel (v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C) \right) \tag{74}$$

where

$$j_1[[\varphi_1]]^m(l_{\text{env}}) \triangleleft q_1]^\bullet \xrightarrow{\tau}^k (v \ \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B) \text{ for } k < n, q_1 < q \tag{75}$$

$$j_2[[\varphi_2]]^m(l_{\text{env}}) \triangleleft q_2]^\bullet \xrightarrow{\tau}^l (v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C) \text{ for } l < n, q_2 < q \tag{76}$$

By Lemma 11, (75) and (76) we also infer that the only external action that can be performed by the processes $(v \ \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B)$ and $(v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C)$ is fail!. Moreover by (75) and (76) we can also show that

$$\text{fId}\left((v \ \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B)\right) = \{j_1\} \quad \text{fId}\left((v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C)\right) = \{j_2\}$$

Thus these two subactors cannot communicate with each other or send messages to actor i . This also means that the remaining possible actions that A can perform are:

$$A \xrightarrow{\tau} (v \ j_1, j_2) \left(i[u, \text{fork}(j_1, j_2) \triangleleft q]^\bullet \parallel (v \ \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B) \parallel (v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2 : u]^\bullet \parallel C) \right) \text{ or} \tag{77}$$

$$A \xrightarrow{\tau} (v \ j_1, j_2) \left(i[j_2!u, \text{fork}(j_1, j_2) \triangleleft q]^\bullet \parallel (v \ \tilde{h}'_1)(j_1[e'_1 \triangleleft q''_1]^\bullet \parallel B') \parallel (v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C) \right)$$

because $(v \ \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B) \xrightarrow{\tau} (v \ \tilde{h}'_1)(j_1[e'_1 \triangleleft q''_1]^\bullet \parallel B')$ or (78)

$$A \xrightarrow{\tau} (v \ j_1, j_2) \left(i[j_2!u, \text{fork}(j_1, j_2) \triangleleft q]^\bullet \parallel (v \ \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B) \parallel (v \ \tilde{h}'_2)(j_2[e'_2 \triangleleft q''_2]^\bullet \parallel C') \right)$$

because $(v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C) \xrightarrow{\tau} (v \ \tilde{h}'_2)(j_2[e'_2 \triangleleft q''_2]^\bullet \parallel C')$ or (79)

$$A \xrightarrow{i?v} (v \ j_1, j_2) \left(i[j_2!u, \text{fork}(j_1, j_2) \triangleleft q : v]^\bullet \parallel (v \ \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^\bullet \parallel B) \parallel (v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2]^\bullet \parallel C) \right) \tag{80}$$

We prove confluence for the pair of actions (77) and (79) and leave the other combinations for the interested reader. From (79) and Lemma 8 we derive

$$(v \ \tilde{h}_2)(j_2[e_2 \triangleleft q'_2 : u]^\bullet \parallel C) \xrightarrow{\tau} (v \ \tilde{h}'_2)(j_2[e'_2 \triangleleft q''_2 : u]^\bullet \parallel C')$$

and by PAR and SCP we obtain

$$\begin{aligned}
 & (\nu j_1, j_2) \left(i[u, \text{fork}(j_1, j_2) \triangleleft q]^{\bullet} \parallel \left(\begin{array}{l} (\nu \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^{\bullet} \parallel B) \\ (\nu \tilde{h}_2)(j_2[e_2 \triangleleft q'_2 : u]^{\bullet} \parallel C) \end{array} \right) \right) \xrightarrow{\tau} \\
 & (\nu j_1, j_2) \left(i[u, \text{fork}(j_1, j_2) \triangleleft q]^{\bullet} \parallel \left(\begin{array}{l} (\nu \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^{\bullet} \parallel B) \\ (\nu \tilde{h}'_2)(j_2[e'_2 \triangleleft q''_2 : u]^{\bullet} \parallel C') \end{array} \right) \right)
 \end{aligned} \tag{81}$$

Using COM, STR, PAR and SCP we can derive

$$\begin{aligned}
 & (\nu j_1, j_2) \left(i[j_2!u, \text{fork}(j_1, j_2) \triangleleft q]^{\bullet} \parallel \left(\begin{array}{l} (\nu \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^{\bullet} \parallel B) \\ (\nu \tilde{h}'_2)(j_2[e'_2 \triangleleft q''_2]^{\bullet} \parallel C') \end{array} \right) \right) \xrightarrow{\tau} \\
 & (\nu j_1, j_2) \left(i[u, \text{fork}(j_1, j_2) \triangleleft q]^{\bullet} \parallel \left(\begin{array}{l} (\nu \tilde{h}_1)(j_1[e_1 \triangleleft q'_1]^{\bullet} \parallel B) \\ (\nu \tilde{h}'_2)(j_2[e'_2 \triangleleft q''_2 : u]^{\bullet} \parallel C') \end{array} \right) \right)
 \end{aligned} \tag{82}$$

thus we close the confluence diamond by (81) and (82). □

Lemma 18 (Weak Confluence) *For all $\varphi \in \text{SHML}$, $q \in \text{VAL}^*$*

$$i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^{\bullet} \implies A \text{ implies } \text{cnf}(A)$$

Proof By strong induction on n , the number of transitions in $i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^{\bullet} (\xrightarrow{\tau})^n A$.

$n = 0$ We know $A = i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^{\bullet}$. It is confluent because it can perform either of two actions, namely a τ -action for the function application (see APP in Fig. 2), or else an external input at i_{mtr} , (see RCVU in Fig. 2). The matching moves can be constructed by RCVU on the one hand, and by Lemma 8 on the other, analogously to the base case of Lemma 17.

$n = k + 1$ By performing an analysis similar to that of Lemma 11, but for $i_{\text{mtr}}[\text{Mon}(\varphi) \triangleleft q]^{\bullet}$ instead, we can determine that this actor can only weakly transition to either of the forms below whereby, for cases (ii) to (v), we obtain B as a result of $i[\llbracket\varphi\rrbracket^{\text{m}}(l_{\text{env}}) \triangleleft r]^{\bullet} \implies B$ for some r :

- (i) $A = i_{\text{mtr}}[M = \text{spw}(\llbracket\varphi\rrbracket^{\text{m}}(\text{nil})), \text{mLoop}(M) \triangleleft q]^{\bullet}$
- (ii) $A \equiv (\nu i)(i_{\text{mtr}}[\text{mLoop}(i) \triangleleft q]^{\bullet} \parallel B)$
- (iii) $A \equiv (\nu i)(i_{\text{mtr}}[\text{rcv } z \rightarrow i!z \text{ end}, \text{mLoop}(i) \triangleleft q]^{\bullet} \parallel B)$
- (iv) $A \equiv (\nu i)(i_{\text{mtr}}[i!v, \text{mLoop}(i) \triangleleft q]^{\bullet} \parallel B)$
- (v) $A \equiv (\nu i)(i_{\text{mtr}}[v, \text{mLoop}(i) \triangleleft q]^{\bullet} \parallel B)$

We here focus on the 4th case of monitor structure; the other cases are analogous. From $i[\llbracket\varphi\rrbracket^{\text{m}}(l_{\text{env}}) \triangleleft r]^{\bullet} \implies B$ and Lemma 11 we know that

$$\begin{aligned}
 B & \xrightarrow{\gamma} \text{ implies } \gamma = \text{fail!} \text{ or } \gamma = \tau \\
 B & \equiv (\nu h)(i[e \triangleleft r]^{\bullet} \parallel C) \text{ where } \text{fid}(B) = i
 \end{aligned}$$

This means that $(\nu i)(i_{\text{mtr}}[i!v, \text{mLoop}(i) \triangleleft q]^{\bullet} \parallel B)$ can only exhibit the following actions:

$$\begin{aligned}
 & (\nu i)(i_{\text{mtr}}[i!v, \text{mLoop}(i) \triangleleft q]^{\bullet} \parallel B) \xrightarrow{i_{\text{mtr}}?u} \\
 & (\nu i)(i_{\text{mtr}}[i!v, \text{mLoop}(i) \triangleleft q : u]^{\bullet} \parallel B)
 \end{aligned} \tag{83}$$

$$\begin{aligned}
 (\nu i)(i_{\text{mtr}}[i!v, \text{mLoop}(i) \triangleleft q]^* \parallel B) &\xrightarrow{\tau} \\
 (\nu i)(i_{\text{mtr}}[v, \text{mLoop}(i) \triangleleft q]^* \parallel (\nu \mathbf{h})(i[e \triangleleft r : v]^\bullet \parallel C)) &\quad (84)
 \end{aligned}$$

$$\begin{aligned}
 (\nu i)(i_{\text{mtr}}[i!v, \text{mLoop}(i) \triangleleft q]^* \parallel B) &\xrightarrow{\tau} (\nu i)(i_{\text{mtr}}[i!v, \text{mLoop}(i) \triangleleft q]^* \parallel B') \\
 &\quad (85)
 \end{aligned}$$

Most pairs of action can be commuted easily by PAR and SCP as they concern distinct elements of the actor system. The only non-trivial case is the pair of actions (84) and (85), which can be commuted using Lemma 8, in analogous fashion to the base case. \square

References

- Aceto L, Ingólfssdóttir A (1999) Testing Hennessy–Milner logic with recursion. In: FoSSaCS'99. Springer, pp 41–55
- Aceto L, Ingólfssdóttir A, Larsen KG, Srba J (2007) Reactive systems: modelling, Specification and verification. Cambridge University Press, New York
- Armstrong J (2007) Programming Erlang. The Pragmatic Bookshelf, Armstrong
- Bauer A, Falcone Y (2012) Decentralised LTL monitoring. In: Giannakopoulou D, Mry D (eds) FM. LNCS, vol 7436. Springer, pp 85–100
- Bauer A, Leucker M, Schallhar C (2011) Runtime verification for LTL and TLTL. ACM Trans Softw Eng Methodol 20:14:1–14:64
- Bocchi L, Chen T-C, Demangeon R, Honda K, Yoshida N (2013) Monitoring networks through multiparty session types. In: FMOODS/FORTE 2013. LNCS, vol 7892. Springer, pp 50–65.
- Cao T-D, Phan-Quang T-T, Felix P, Castanet R (2010) Automated runtime verification for web services. In: ICWS. IEEE, pp 76–82
- Carlsson R (2001) An introduction to Core Erlang. In: PLI'01 (Erlang Workshop)
- Cassar I, Francalanza A (2014) On synchronous and asynchronous monitor instrumentation for actor-based systems. In: FOCLASA, EPTCS (to appear)
- Cerone A, Hennessy M (2010) Process behaviour: formulae vs. tests. In: EXPRESS'10, vol 41 EPTCS, pp 31–45
- Cesarini F, Thompson S (2009) Erlang programming. O'Reilly, Sebastopol
- Chang E, Manna Z, Pnueli A (1992) Characterization of temporal property classes. In: ALP. LNCS, vol 623. Springer-Verlag, pp 474–486
- Clarke E Jr, Grumberg O, Peled D (1999) Model checking. MIT Press, Cambridge
- Colombo C, Francalanza A, Gatt R (2011) Elarva: a monitoring tool for Erlang. In: RV. LNCS, vol 7186. Springer, pp 370–374
- Colombo C, Francalanza A, Grima I (2012) Simplifying contract-violating traces. In: FLACOS, EPTCS, vol 94, pp 11–20
- Colombo C, Francalanza A, Mizzi R, Pace GJ (2012) polylarva: runtime verification with configurable resource-aware monitoring boundaries. In: SEFM, pp 218–232
- D'Angelo B, Sankaranarayanan S, Sánchez C, Robinson W, Finkbeiner B, Sipma HB, Mehrotra S, Manna Z. (2005) Lola: runtime monitoring of synchronous systems. In: TIME, IEEE
- Falcone Y, Jaber M, Nguyen T-H, Bozga M, Bensalem S. (2011) Runtime verification of component-based systems. In: SEFM. LNCS, vol 7041. Springer, pp 204–220
- Francalanza A, Seychell A (2013) Synthesising correct concurrent runtime monitors in Erlang. Technical Report CS2013-01, University of Malta. <https://www.cs.um.edu.mt/svrg/papers.html>. Accessed Jan
- Francalanza A, Gauci A, Pace GJ (2013) Distributed System contract monitoring. JLAP 82(5–7):186–215
- Francalanza A, Seychell A, Cassar I. DetectEr. <https://bitbucket.org/casian/detector2.0>
- Fredlund L-Å (2001) A framework for reasoning about Erlang code. PhD thesis, Royal Institute of Technology, Stockholm, Sweden
- Geilen M (2001) On the construction of monitors for temporal logic properties. ENTCS 55(2):181–199
- Hennessy M (2008) A distributed picalculus. Cambridge University Proess, Cambridge
- Hennessy M, Milner R (1985) Algebraic laws for nondeterminism and concurrency. J ACM 32(1): 137–161
- Hewitt C, Bishop P, Steiger R (1973) A universal modular actor formalism for artificial intelligence. In: IJCAI, Morgan Kaufmann, pp 235–245
- Kozen D (1983) Results on the propositional μ -calculus. TCS 27:333–354

28. Manna Z, Pnueli A (1990) A hierarchy of temporal properties (invited paper, 1989). In: PODC, ACM, pp 377–410
29. Meredith PO, Jin D, Griffith D, Chen F, Rosu G (2012) An overview of the MOP runtime verification framework. *STTT* 14(3):249–289
30. Milner R (1989) *Communication and concurrency*. Prentice-Hall Inc, Upper Saddle River
31. Milner R, Parrow J, Walker D (1993) Modal logics for mobile processes. *TCS* 114:149–171
32. Nicola RD, Hennessy MCB (1984) Testing equivalences for processes, *TCS*, pp 83–133
33. Rensink A, Vogler W (2007) Fair testing. *Inf Comput* 205(2):125–198
34. Sen K, Rosu G, Agha G (2004) Generating optimal linear temporal logic monitors by coinduction. In: *ASIAN. LNCS*, vol 2896. Springer-Verlag, pp 260–275
35. Sen K, Vardhan A, Agha G, Roşu G (2004) Efficient decentralized monitoring of safety in distributed systems. *ICSE*, pp 418–427
36. Svensson H, Fredlund L-Å, Benac Earle C (2010) A unified semantics for future erlang. In: *Erlang Workshop*, ACM, pp 23–32