

On Verifying Hennessy-Milner Logic with Recursion at Runtime^{*}

Adrian Francalanza¹, Luca Aceto², and Anna Ingólfssdóttir²

¹ CS, ICT, University of Malta, Malta `adrian.francalanza@um.edu.mt`

² ICE-TCS, Reykjavik University, Iceland `{luca,anna}@ru.is`

Abstract. We study μ HML (a branching-time logic with least and greatest fixpoints) from a runtime verification perspective. We establish which subset of the logic can be verified at runtime and define correct monitor-synthesis algorithms for this subset. We also prove completeness results *wrt.* these logical subsets that show that no other properties apart from those identified can be verified at runtime.

1 Introduction

Runtime Verification (RV) [20, 14] is a lightweight verification technique whereby the *execution of a system* is analysed with the aim of inferring correctness *wrt.* some property. Despite its advantages, the technique is generally limited when compared to other verification techniques such as model-checking because certain correctness properties cannot be verified at runtime [21, 10]. For instance, *online* RV analyses *partial* executions incrementally (up to the current execution point of the system) which limits its applicability to satisfaction verdicts relating to correctness properties describing complete (*i.e.*, potentially infinite) executions.

There are broadly two approaches to address such a limitation. The first approach is to *restrict* the expressive power of the correctness specifications: typically, one either limits specifications to *descriptions of finite traces* such as regular expressions (RE) [12, 15], or else *redefines* the semantics of existing logics (*e.g.*, LTL) so as to reflect the limitations of the runtime setting [13, 5, 7, 6]. The second approach is to leave the semantics of the specification logic *unchanged*, and study which *subsets* of the logic can be verified at runtime [17, 24, 11, 16].

Both approaches have their merits. The first approach is, in general, more popular and tends to produce specifications that are closely related to the monitors that check for them (*e.g.*, RE and automata in [15, 23]), thus facilitating aspects such as monitor correctness. On the other hand, the second approach does not hinder the expressive power of the logic. Instead, it allows a verification framework to determine whether either to check for a property at runtime (when possible), or else to employ more powerful (and expensive) verification techniques such as model-checking. One can even envisage a hybrid approach,

^{*} The research of L. Aceto and A. Ingólfssdóttir was supported by the project 001-ABEL-CM-2013 of the NILS Science and Sustainability Programme and the project Nominal SOS (nr. 141558-051) of the Icelandic Research Fund.

Syntax

$$p, q, r \in \text{PROC} ::= \text{nil} \quad (\text{inaction}) \quad | \quad \alpha.p \quad (\text{prefixing}) \quad | \quad p + q \quad (\text{choice}) \\ | \quad \text{rec } x.p \quad (\text{recursion}) \quad | \quad x \quad (\text{rec. variable})$$

Dynamics

$$\text{ACT} \frac{}{\alpha.p \xrightarrow{\alpha} p} \quad \text{REC} \frac{}{\text{rec } x.p \xrightarrow{\tau} p[\text{rec } x.p/x]} \quad \text{SELL} \frac{p \xrightarrow{\mu} p'}{p + q \xrightarrow{\mu} p'}$$

Fig. 1. A Model for describing Systems

where parts of a property are verified using RV and other parts are checked using other techniques. More importantly, however, the second approach leads to better separation of concerns: since it is *agnostic* of the verification technique used, one can change the method of verification without impinging on the property semantics.

This paper follows this second approach. In particular, it revisits Hennessy-Milner Logic with recursion [3], μHML , a reformulation of the expressive modal μ -calculus [19], used to describe correctness properties of reactive system; subsets of the logic have already been adapted for `detectEr` [25], an RV tool for runtime-verifying actor-based systems [16, 8], whereas constructs from the modal μ -calculus have been used in other RV tools such as `Eagle` [4]. In this study we consider the logic in its entirety, and investigate the monitorability of the logic *wrt.* an operational definition of a general class of monitors that employ both acceptance and rejection verdicts [7, 14]. In particular, our results extend the class of monitorable μHML properties used in [16] and establish monitorability upper bounds for this logic. We also present new results that relate the utility of multi-verdict monitors *wrt.* logics defined over programs (as opposed to traces). To the best of our knowledge, this is one of the first bodies of work investigating the limits of RV *wrt.* a *branching-time* logic that specifies properties about the *execution graph* of a program; other work pertaining to the aforementioned second approach has focussed on *linear-time* logics defined over *execution traces*, and has explored RV's limits along the linear-time dimension, *e.g.*, [11].

In the rest of the paper, Sec. 2 introduces our model for reactive systems and Sec. 3 presents the logic μHML defined over this model. Sec. 4 formalises our abstract RV operational setup in terms of monitors and our instrumentation relation. In Sec. 5 we argue for a particular correspondence between monitors and μHML properties within this setup. Sec. 6 identifies monitorability limits for the logic but also establishes a monitorable logical subset that satisfies the correspondence of Sec. 5. Sec. 7 shows that this subset is maximally expressive, using a result about multi-verdict monitors. Sec. 8 concludes.




Fig. 2. A depiction of the system in Example 1

2 The Model

We describe systems abstractly as Labelled Transition Systems (LTSs) [3, 19]. An LTS is a triple $\langle \text{PROC}, (\text{ACT} \cup \{\tau\}), \longrightarrow \rangle$ consisting of a set of states, PROC, a set of actions, ACT with distinguished silent action τ (we assume $\mu \in \text{ACT} \cup \{\tau\}$ and $\tau \notin \text{ACT}$), and a transition relation, $\longrightarrow \subseteq (\text{PROC} \times (\text{ACT} \cup \{\tau\}) \times \text{PROC})$. LTS states can be expressed as processes, PROC, from the regular fragment of CCS [22] as defined in Fig. 1. Assuming a set of (visible) actions, $\alpha, \beta \in \text{ACT}$ and a set of (recursion) variables $x, y, z \in \text{VARS}$, processes may be either inactive, prefixed by an action, a mutually-exclusive choice amongst two processes, or recursive; $\text{rec } x.p$ acts as a *binder* for x in p and we work up to alpha-conversion of bound variables. All recursive processes are assumed to be guarded.

The dynamic behaviour is then described by the transition rules of Fig. 1, defined over the closed and guarded terms in PROC (we elide the symmetric rule SELR). The suggestive notation $p \xrightarrow{\mu} p'$ denotes $(p, \mu, p') \in \longrightarrow$; we also write $p \not\xrightarrow{\alpha}$ to denote $\neg(\exists p'. p \xrightarrow{\alpha} p')$. For example, $p_1 + p_2 \xrightarrow{\mu} q$ if either $p_1 \xrightarrow{\mu} q$ or $p_2 \xrightarrow{\mu} q$. As usual, we write $p \Longrightarrow p'$ in lieu of $p(\tau \rightarrow)^* p'$ and $p \xRightarrow{\mu} p'$ for $p \Longrightarrow \cdot \xrightarrow{\mu} \cdot \Longrightarrow p'$, referring to p' as a μ -derivative of p . We let $t, u \in \text{ACT}^*$ range over sequences of visible actions and write $p \xRightarrow{\alpha_1} \dots \xRightarrow{\alpha_n} p_n$ as $p \xRightarrow{t} p_n$, where $t = \alpha_1, \dots, \alpha_n$. See [22, 3] for more details.

Example 1. A (reactive) system that acts as a server that repeatedly accepts requests and subsequently answers them, with the possibility of terminating through the special close request, may be expressed as the following process, p .

$$p = \text{rec } x. (\text{req.ans}.x + \text{cls.nil})$$

Syntax

$\varphi, \phi \in \mu\text{HML} ::= \text{tt}$	(truth)	ff	(falsehood)
$\varphi \vee \phi$	(disjunction)	$\varphi \wedge \phi$	(conjunction)
$\langle \alpha \rangle \varphi$	(possibility)	$[\alpha] \varphi$	(necessity)
$\min X. \varphi$	(min. fixpoint)	$\max X. \varphi$	(max. fixpoint)
X	(rec. variable)		

Semantics

$\llbracket \text{tt}, \rho \rrbracket \stackrel{\text{def}}{=} \text{PROC}$	$\llbracket \text{ff}, \rho \rrbracket \stackrel{\text{def}}{=} \emptyset$
$\llbracket \varphi_1 \wedge \varphi_2, \rho \rrbracket \stackrel{\text{def}}{=} \llbracket \varphi_1, \rho \rrbracket \cap \llbracket \varphi_2, \rho \rrbracket$	$\llbracket \varphi_1 \vee \varphi_2, \rho \rrbracket \stackrel{\text{def}}{=} \llbracket \varphi_1, \rho \rrbracket \cup \llbracket \varphi_2, \rho \rrbracket$
$\llbracket [\alpha] \varphi, \rho \rrbracket \stackrel{\text{def}}{=} \{p \mid p \xrightarrow{\alpha} q \text{ implies } q \in \llbracket \varphi, \rho \rrbracket\}$	$\llbracket \langle \alpha \rangle \varphi, \rho \rrbracket \stackrel{\text{def}}{=} \{p \mid p \xrightarrow{\alpha} q \text{ and } q \in \llbracket \varphi, \rho \rrbracket\}$
$\llbracket \min X. \varphi, \rho \rrbracket \stackrel{\text{def}}{=} \bigcap \{S \mid \llbracket \varphi, \rho[X \mapsto S] \rrbracket \subseteq S\}$	
$\llbracket \max X. \varphi, \rho \rrbracket \stackrel{\text{def}}{=} \bigcup \{S \mid S \subseteq \llbracket \varphi, \rho[X \mapsto S] \rrbracket\}$	$\llbracket X, \rho \rrbracket \stackrel{\text{def}}{=} \rho(X)$

Fig. 3. μHML Syntax and Semantics

A server that non-deterministically stops offering the close action is denoted by process q , whereas r only offers the close action after the first serviced request.

$$q = \text{rec } x. (\text{req.ans.}x + \text{cls.nil} + (\text{rec } y. \text{req.ans.}y))$$

$$r = \text{req.ans.rec } x. (\text{req.ans.}x + \text{cls.nil})$$

Pictorially, the *resp.* LTSs denoted by processes p, q and r are shown in Fig. 2, where the arcs correspond to weak transitions, $\xRightarrow{\mu}$. ■

3 The Logic

The logic μHML assumes a countable set of logical variables $X, Y \in \text{LVAR}$, and is defined as the set of *closed* formulae generated by the grammar of Fig. 3. Apart from the standard constructs for truth, falsehood, conjunction and disjunction, the logic is equipped with possibility and necessity modal operators, together with recursive formulae expressing least or greatest fixpoints; formulae $\min X. \varphi$ and $\max X. \varphi$ *resp.* bind free instances of the logical variable X in φ , inducing the usual notions of open/closed formulae and equality up to alpha-conversion.

Formulae are interpreted over the process powerset domain, $S \in \mathcal{P}(\text{PROC})$. The semantic definition of Fig. 3 is given for *both* open and closed formulae and employs an environment from variables to sets of processes, $\rho \in \text{LVAR} \rightarrow \mathcal{P}(\text{PROC})$; this permits an inductive definition on the structure of the formula. For instance, in Fig. 3, the semantic meaning of a variable X wrt. an environment ρ is the mapping for that variable in ρ . The semantics of truth, falsehood, conjunction and disjunction are standard (*e.g.*, \vee and \wedge are interpreted as set-theoretic union and intersection). Possibility formulae $\langle \alpha \rangle \varphi$ describe processes

with *at least one* α -derivative satisfying φ whereas necessity formulae $[\alpha]\varphi$ describe processes where *all* of their α -derivatives (possibly none) satisfy φ . The powerset domain $\mathcal{P}(\text{PROC})$ is a complete lattice wrt. set-inclusion, \subseteq , which guarantees the existence of least and largest solutions for the recursive formulae of the logic; as usual, these can be *resp.* specified as the intersection of all the pre-fixpoint solutions and the union of all post-fixpoint solutions [3]. Note that $\rho[X \mapsto S]$ denotes an environment ρ' where $\rho'(X) = S$ and $\rho'(Y) = \rho(Y)$ for all other $Y \neq X$. Since the interpretation of closed formulae is independent of the environment ρ , we sometimes write $\llbracket \varphi \rrbracket$ in lieu of $\llbracket \varphi, \rho \rrbracket$. We say that a process p *satisfies* a formula φ whenever $p \in \llbracket \varphi \rrbracket$, and *violates* a formula whenever $p \notin \llbracket \varphi \rrbracket$.

Example 2. Formula $\langle \alpha \rangle \text{tt}$ describes processes that *can* perform action α whereas formula $[\alpha] \text{ff}$ describes processes that *cannot* perform action α .

$$\begin{aligned} \varphi_1 &= \min X.(\langle \text{req} \rangle \langle \text{ans} \rangle X \vee [\text{cls}] \text{ff}) & \varphi_2 &= \max X.(\langle \text{req} \rangle \langle \text{ans} \rangle X \vee [\text{cls}] \text{ff}) \\ \varphi_3 &= \max X.([\text{req}] [\text{ans}] X \wedge \langle \text{cls} \rangle \text{tt}) & \varphi_4 &= \max X.([\text{req}] [\text{ans}] X \wedge [\text{cls}] \text{ff}) \end{aligned}$$

Formula φ_1 denotes a *liveness* property describing processes that *eventually* stop offering the action `cls` after any number of serviced request, $(\text{req.ans})^*$ —processes q and r from Ex. 1 satisfy this property but p does not. Changing the fixpoint into a maximal one, *i.e.*, φ_2 , would include p in the property as well. Formulae φ_3 and φ_4 denote *safety* properties: *e.g.*, φ_3 , describes (terminating and non-terminating) processes that can *always* perform a `cls` action after any number of serviced request (p satisfies this property but q and r do not).

$$\begin{aligned} \varphi_5 &= \langle \text{req} \rangle \langle \text{ans} \rangle \max X.([\text{req}] \text{ff} \vee \langle \text{req} \rangle \langle \text{ans} \rangle X) \wedge [\text{cls}] \text{ff} \\ \varphi_6 &= \min X.(\langle \text{req} \rangle \langle \text{ans} \rangle \text{tt} \wedge [\text{req}] [\text{ans}] X) \vee \langle \text{cls} \rangle \text{tt} \end{aligned}$$

Formula φ_5 is satisfied by processes that, after one serviced request, `req.ans`, exhibit a complete³ transition sequence $(\text{req.ans})^*$ that never offers action `cls` (q from Ex. 1 satisfies this property whereas p and r do not). Formula φ_6 describes processes that along *all* serviced request sequences, $(\text{req.ans})^*$, eventually reach a stage where they offer action `cls` (processes p and q satisfy the criteria immediately, whereas r satisfies it for $(\text{req.ans})^*$ sequences longer than 1). ■

4 Monitors and instrumentation

Monitors may also be viewed as LTSs, through the syntax of Fig. 4; this is similar to that of processes, with the exception that `nil` is replaced by three *verdict* constructs, `yes`, `no` and `end`, *resp.* denoting acceptance, rejection and termination (*i.e.*, an inconclusive outcome). Monitor behaviour is similar to that of processes for the common constructs; see rules in Fig. 4. The only new transition concerns verdicts, `mVER`, stating that a verdict may transition with *any* $\alpha \in \text{ACT}$ and go back to the same state, modelling the requirement that verdicts are *irrevocable*.

³ A transition sequence is complete if it is either infinite or affords no more actions.

Syntax

$$\begin{array}{l}
m, n \in \text{MON} ::= v \quad | \quad \alpha.m \quad | \quad m + n \quad | \quad \text{rec } x.m \quad | \quad x \\
v, u \in \text{VERD} ::= \text{end} \quad | \quad \text{no} \quad | \quad \text{yes}
\end{array}$$

Dynamics

$$\begin{array}{l}
\text{MACT} \frac{}{\alpha.m \xrightarrow{\alpha} m} \quad \text{MREC} \frac{}{\text{rec } x.m \xrightarrow{\tau} m[\text{rec } x.m/x]} \\
\text{MSELL} \frac{m \xrightarrow{\mu} m'}{m + n \xrightarrow{\mu} m'} \quad \text{MSELR} \frac{n \xrightarrow{\mu} n'}{m + n \xrightarrow{\mu} n'} \quad \text{MVER} \frac{}{v \xrightarrow{\alpha} v}
\end{array}$$

Instrumentation

$$\begin{array}{l}
\text{IMON} \frac{p \xrightarrow{\alpha} p' \quad m \xrightarrow{\alpha} m'}{m \triangleleft p \xrightarrow{\alpha} m' \triangleleft p'} \quad \text{ITER} \frac{p \xrightarrow{\alpha} p' \quad m \not\xrightarrow{\alpha} \quad m \not\xrightarrow{\tau}}{m \triangleleft p \xrightarrow{\alpha} \text{end} \triangleleft p'} \\
\text{IASYP} \frac{p \xrightarrow{\tau} p'}{m \triangleleft p \xrightarrow{\tau} m \triangleleft p'} \quad \text{IASYM} \frac{m \xrightarrow{\tau} m'}{m \triangleleft p \xrightarrow{\tau} m' \triangleleft p}
\end{array}$$

Fig. 4. Monitors and Instrumentation

Fig. 4 also describes an instrumentation relation connecting the behaviour of a process p with that of a monitor m : the configuration $m \triangleleft p$ denotes a *monitored system*. In an instrumentation, the process leads the (visible) behaviour of a monitored system (*i.e.*, if the process cannot α -transition, then the monitored system will not either) while the monitor passively follows, transitioning accordingly; this is in contrast with well-studied parallel composition relations of LTSs [22, 18]. Specifically, rule **IMON** states that if a process can transition with action α and the *resp.* monitor can follow this by transitioning with the same action, then in an instrumented monitored system they transition in lockstep. However, if the monitor cannot follow such a transition, $m \not\xrightarrow{\alpha}$, even after any number of internal actions, $m \not\xrightarrow{\tau}$, instrumentation forces it to terminate with an inconclusive verdict, **end**, while the process is allowed to proceed unaffected; see rule **ITER**. Rules **IASYP** and **IASYM** allow monitors and processes to transition independently *wrt.* internal moves.⁴

Proposition 1. $m \triangleleft p \xrightarrow{t} m' \triangleleft p' \quad \text{iff} \quad p \xrightarrow{t} p' \text{ and}$

- either $m \xrightarrow{t} m'$
- or $m' = \text{end}$ and $\exists t', \alpha, t'', m''$. $t = t'\alpha t''$, $m'' \not\xrightarrow{\alpha}$ and $m \xrightarrow{t'} m'' \not\xrightarrow{\alpha}$.

Remark 1. Since we strive towards a general theory, the syntax in Fig. 4 allows for non-deterministic monitors such as $\alpha.\text{yes} + \alpha.\text{no}$ or $\alpha.\text{nil} + \alpha.\beta.\text{yes}$. There are

⁴ If a monitor cannot match a process action, but can transition silently, it is allowed to do so, and the matching check is applied again to the τ -derivative monitor.

settings where determinism is unattainable (e.g., distributed monitoring [15]) or desirable (e.g., testers [23]), and others where non-determinism expresses under-specification (e.g., program refinement [1]). Thus, expressing non-determinism allows us to study the cases where it is tolerated or considered erroneous.

Example 3. Monitor m_1 (defined below) monitors for executions of the form $(\text{req.ans})^*.\text{cls}$ returning the acceptance verdict **yes**, whereas m_2 dually rejects executions of that form. When composed with process p from Ex. 1, the monitored system $m_1 \triangleleft p$ may either service requests forever, $m_1 \triangleleft p \xrightarrow{\text{req}} \cdot \xrightarrow{\text{ans}} m_1 \triangleleft p$, or else terminate with a **yes** verdict, $m_1 \triangleleft p \xrightarrow{\text{cls}} \text{yes} \triangleleft \text{nil}$. By contrast, when instrumented over a process capable of the transition $p' \xrightarrow{\text{ans}} p''$, m_1 may terminate its process observation after one transition, i.e., $m_1 \triangleleft p' \xrightarrow{\text{ans}} \text{end} \triangleleft p''$.

$$\begin{aligned} m_1 &= \text{rec } x.(\text{req.ans}.x + \text{cls.yes}) & m_3 &= \text{rec } x.(\text{req.ans}.x + \text{cls.yes} + \text{req.req}.x) \\ m_2 &= \text{rec } x.(\text{req.ans}.x + \text{cls.no}) & m_4 &= \text{rec } x.(\text{req.ans}.x + \text{cls.yes} + \text{cls.no}) \end{aligned}$$

Monitor m_3 may either behave like m_1 or non-deterministically terminate upon a serviced request, i.e., $m_3 \triangleleft p \xrightarrow{\text{req}} \text{req}.m_3 \triangleleft p_1 \xrightarrow{\text{ans}} \text{end} \triangleleft p$. Conversely, monitor m_4 non-deterministically returns verdict **yes** or **no** upon a **cls** action, e.g., $m_4 \triangleleft p \xrightarrow{\text{cls}} \text{yes} \triangleleft \text{nil}$ but also $m_4 \triangleleft p \xrightarrow{\text{cls}} \text{no} \triangleleft \text{nil}$. ■

5 Correspondence

Our goal is to establish a correspondence between the verdicts reached by monitors over an instrumented system from Sec. 4 and the properties specified using the logic of Sec. 3. In particular, we would like to relate acceptances (**yes**) and rejections (**no**) reached by a monitor m when monitoring a process p with satisfactions ($p \in \llbracket \varphi \rrbracket$) and violations ($p \notin \llbracket \varphi \rrbracket$) for that process wrt. some μHML formula, φ . This will, in turn, allow us to determine when a monitor m *represents* (in some precise sense) a property φ .

Example 4. Monitor m_1 from Ex. 3 monitors for *satisfactions* of the property

$$\varphi_7 = \min X.(\langle \text{req} \rangle \langle \text{ans} \rangle X \vee \langle \text{cls} \rangle \text{tt})$$

describing processes that can perform a **cls** action after a number of serviced requests. Stated otherwise, m_1 produces a **yes** verdict for a computation of the form $(\text{req.ans})^*.\text{cls}$ from a process p , attesting that $p \in \llbracket \varphi_7 \rrbracket$. Similarly, m_2 from Ex. 3 monitors for *violations* of the property φ_4 from Ex. 2. The same *cannot* be said for m_4 from Ex. 3 and φ_7 above: for some processes, e.g., p from Ex. 1, it may produce both verdicts **yes** and **no** for witness computations $(\text{req.ans})^*.\text{cls}$, which leads to contradictions at a logical level i.e., we cannot have both $p \in \llbracket \varphi_7 \rrbracket$ and $p \notin \llbracket \varphi_7 \rrbracket$. A similar argument applies to m_4 and m_2 from Ex. 3. ■

Remark 2. A monitor may behave non-deterministically in other ways wrt. a process. For instance, m_3 from Ex. 3 may sometimes flag an acceptance but

at other times may not when monitoring p from Ex. 1, *even when* p produces the same trace: e.g., for $t = \text{req.ans.cls}$ we have $m_3 \triangleleft p \xrightarrow{t} \text{yes} \triangleleft \text{nil}$ but also $m_3 \triangleleft p \xrightarrow{t} \text{end} \triangleleft \text{nil}$. However, since any other terminal monitor state apart from **yes** and **no** (i.e., **end** and any other non-verdict state) is of no consequence from a logical satisfaction/violation point of view, we abstract from such outcomes.

We investigate conditions that one could require for establishing the correspondence between monitors and formulae. We start with Def. 2 (below): it defines when m is able to *monitor soundly for a property* φ , $\mathbf{smon}(m, \varphi)$, by requiring that acceptances (*resp.* rejections) imply satisfactions (*resp.* violations) for every monitored execution of a process p .

Definition 1 (Acceptance/Rejection). $\mathbf{acc}(p, m) \stackrel{\text{def}}{=} \exists t, p'. m \triangleleft p \xrightarrow{t} \text{yes} \triangleleft p'$ and $\mathbf{rej}(p, m) \stackrel{\text{def}}{=} \exists t, p'. m \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$.

Definition 2 (Sound Monitoring).

$\mathbf{smon}(m, \varphi) \stackrel{\text{def}}{=} \forall p. (\mathbf{acc}(p, m) \text{ implies } p \in \llbracket \varphi \rrbracket) \text{ and } (\mathbf{rej}(p, m) \text{ implies } p \notin \llbracket \varphi \rrbracket)$

Note that, if $\mathbf{smon}(m, \varphi)$ and $\exists p. \mathbf{acc}(p, m)$, by Def. 2 we know $p \in \llbracket \varphi \rrbracket$; thus, $\neg(p \notin \llbracket \varphi \rrbracket)$ and by the contrapositive of Def. 2, we must also have $\neg \mathbf{rej}(p, m)$.

Example 5. From Ex. 4, we formally have $\mathbf{smon}(m_1, \varphi_7)$, $\mathbf{smon}(m_2, \varphi_2)$ and $\mathbf{smon}(m_3, \varphi_7)$. We can also show that $\neg \mathbf{smon}(m_4, \varphi_7)$ and $\neg \mathbf{smon}(m_4, \varphi_2)$. ■

Sound monitoring is arguably the least requirement for relating a monitor with a logical property. Further to this, the obvious additional requirement would be to ask for the dual of Def. 2, i.e., *complete monitoring* for m and φ , stating that for all p , $p \in \llbracket \varphi \rrbracket$ implies $\mathbf{acc}(p, m)$, and also that $p \notin \llbracket \varphi \rrbracket$ implies $\mathbf{rej}(p, m)$. However, such a requirement turns out to be too strong for a large part of the logic presented in Fig. 3.

Example 6. Consider the basic formula $\langle \alpha \rangle \text{tt}$. One could ascertain that the simple monitor $\alpha.\text{yes}$ satisfies the condition that $p \in \llbracket \varphi \rrbracket$ implies $\mathbf{acc}(p, m)$ for all p . However, there does not exist a sound monitor that can satisfy $\forall p. p \notin \llbracket \varphi \rrbracket$ implies $\mathbf{rej}(p, m)$ for $\langle \alpha \rangle \text{tt}$. Arguing by contradiction, assume that one such monitor m exists. Since $\text{nil} \notin \llbracket \langle \alpha \rangle \text{tt} \rrbracket$ then we should have $\mathbf{rej}(\text{nil}, m)$. By Def. 1 and Prop. 1, this means $m \Rightarrow \text{no}$ which, in turn, implies that $\mathbf{rej}(\alpha.\text{nil}, m)$ although, clearly, $\alpha.\text{nil} \in \llbracket \langle \alpha \rangle \text{tt} \rrbracket$. This makes m unsound, contradicting our initial assumption.

A similar, albeit dual, argument can be carried out for another core basic formula, $[\alpha] \text{ff}$: although there are sound monitors satisfying the condition $\forall p. p \notin \llbracket \varphi \rrbracket$ implies $\mathbf{rej}(p, m)$, there are none that also satisfy the other condition $\forall p. p \in \llbracket \varphi \rrbracket$ implies $\mathbf{acc}(p, m)$. ■

Concretely, requiring complete monitoring would limit correspondence to a trivial subset of the logic, namely **tt** and **ff**. We therefore define the weaker forms of completeness that are stated below.

Definition 3 (Satisfaction/Violation/Partially-Complete Monitoring).

$$\begin{aligned}
\mathbf{scmon}(m, \varphi) &\stackrel{\text{def}}{=} \forall p.p \in \llbracket \varphi \rrbracket \text{ implies } \mathbf{acc}(p, m) && (\text{satisfaction complete}) \\
\mathbf{vcmon}(m, \varphi) &\stackrel{\text{def}}{=} \forall p.p \notin \llbracket \varphi \rrbracket \text{ implies } \mathbf{rej}(p, m) && (\text{violation complete}) \\
\mathbf{cmon}(m, \varphi) &\stackrel{\text{def}}{=} \mathbf{scmon}(m, \varphi) \text{ or } \mathbf{vcmon}(m, \varphi) && (\text{partially complete})
\end{aligned}$$

We can now formalise monitor-formula correspondence: m monitors for φ , $\mathbf{mon}(m, \varphi)$, if it can do it *soundly*, and in a *partially-complete* manner, i.e., if it is *either* satisfaction complete or violation complete.

Definition 4 (Monitoring). $\mathbf{mon}(m, \varphi) \stackrel{\text{def}}{=} \mathbf{scmon}(m, \varphi) \text{ and } \mathbf{vcmon}(m, \varphi)$.

6 Monitorability

Using Def. 4, we can define what it means for a formula to be monitorable.

Definition 5 (Monitorability). Formula φ is monitorable iff $\exists m.\mathbf{mon}(m, \varphi)$. A language $\mathcal{L} \subseteq \mu\text{HML}$ is monitorable iff every $\varphi \in \mathcal{L}$ is monitorable.

We immediately note that not all logical formulae are monitorable.

Example 7. Through the witness outlined in Ex. 6, we can show that formulae $\langle \alpha \rangle \mathbf{tt}$ and $\langle \beta \rangle \mathbf{tt}$ are monitorable with a satisfaction complete monitor. However, $\varphi_8 = \langle \alpha \rangle \mathbf{tt} \wedge \langle \beta \rangle \mathbf{tt}$ (their conjunction), is *not*. Intuitively, this is so because once a monitor observes one of the actions, it cannot “go back” to check for the other. Formally, we argue towards a contradiction by assuming that $\exists m.\mathbf{mon}(m, \varphi_8)$. There are two subcases to consider.

If m is satisfaction complete, then $\mathbf{acc}(\alpha.\mathbf{nil} + \beta.\mathbf{nil}, m)$ since $\alpha.\mathbf{nil} + \beta.\mathbf{nil} \in \llbracket \varphi_8 \rrbracket$. By Prop. 1, m reaches verdict **yes** along one of the traces ϵ, α or β . If the trace is ϵ , then m also accepts **nil**, which is unsound (since $\mathbf{nil} \notin \llbracket \varphi_8 \rrbracket$) whereas if the trace is α , m must also accept $\alpha.\mathbf{nil}$, which is also unsound ($\alpha.\mathbf{nil} \notin \llbracket \varphi_8 \rrbracket$); the case for β is analogous.

If m is violation complete then $\mathbf{rej}(\beta.\mathbf{nil}, m)$ since $\beta.\mathbf{nil} \notin \llbracket \varphi_8 \rrbracket$. By Prop. 1, we either have $m \xrightarrow{\epsilon} \mathbf{no}$ or $m \xrightarrow{\beta} \mathbf{no}$ and for both cases we can argue that m also rejects process $\alpha.\mathbf{nil} + \beta.\mathbf{nil}$, which is unsound since $\alpha.\mathbf{nil} + \beta.\mathbf{nil} \in \llbracket \varphi_8 \rrbracket$. ■

We now identify a syntactic subset of μHML formulae called mHML , with the aim of showing that it is a monitorable subset of the logic. At an intuitive level, it consists of the safe and co-safe syntactic subsets of μHML , sHML and cHML respectively.

Definition 6 (Monitorable Logic). $\psi, \chi \in \text{mHML} \stackrel{\text{def}}{=} \text{sHML} \cup \text{cHML}$ where:

$$\begin{array}{llllll}
\theta, \vartheta \in \text{sHML} ::= \mathbf{tt} & | \mathbf{ff} & | [\alpha]\theta & | \theta \wedge \vartheta & | \max X.\theta & | X \\
\pi, \varpi \in \text{cHML} ::= \mathbf{tt} & | \mathbf{ff} & | \langle \alpha \rangle \pi & | \pi \vee \varpi & | \min X.\pi & | X
\end{array}$$

To prove monitorability for MHML, we define a monitor synthesis function $\langle - \rangle$ generating a monitor for each $\psi \in \text{MHML}$. We then show that $\langle \psi \rangle$ is the witness monitor required by Def. 5 to demonstrate the monitorability of ψ .

Definition 7 (Monitor Synthesis).

$$\begin{array}{l}
\langle \text{ff} \rangle \stackrel{\text{def}}{=} \text{no} \qquad \qquad \langle \text{tt} \rangle \stackrel{\text{def}}{=} \text{yes} \qquad \qquad \langle X \rangle \stackrel{\text{def}}{=} x \\
\langle [\alpha] \psi \rangle \stackrel{\text{def}}{=} \begin{cases} \alpha. \langle \psi \rangle & \text{if } \langle \psi \rangle \neq \text{yes} \\ \text{yes} & \text{otherwise} \end{cases} \qquad \langle \langle \alpha \rangle \psi \rangle \stackrel{\text{def}}{=} \begin{cases} \alpha. \langle \psi \rangle & \text{if } \langle \psi \rangle \neq \text{no} \\ \text{no} & \text{otherwise} \end{cases} \\
\langle \psi_1 \wedge \psi_2 \rangle \stackrel{\text{def}}{=} \begin{cases} \langle \psi_1 \rangle & \text{if } \langle \psi_2 \rangle = \text{yes} \\ \langle \psi_2 \rangle & \text{if } \langle \psi_1 \rangle = \text{yes} \\ \langle \psi_1 \rangle + \langle \psi_2 \rangle & \text{otherwise} \end{cases} \qquad \langle \psi_1 \vee \psi_2 \rangle \stackrel{\text{def}}{=} \begin{cases} \langle \psi_1 \rangle & \text{if } \langle \psi_2 \rangle = \text{no} \\ \langle \psi_2 \rangle & \text{if } \langle \psi_1 \rangle = \text{no} \\ \langle \psi_1 \rangle + \langle \psi_2 \rangle & \text{otherwise} \end{cases} \\
\langle \max X. \psi \rangle \stackrel{\text{def}}{=} \begin{cases} \text{rec } x. \langle \psi \rangle & \text{if } \langle \psi \rangle \neq \text{yes} \\ \text{yes} & \text{otherwise} \end{cases} \qquad \langle \min X. \psi \rangle \stackrel{\text{def}}{=} \begin{cases} \text{rec } x. \langle \psi \rangle & \text{if } \langle \psi \rangle \neq \text{no} \\ \text{no} & \text{otherwise} \end{cases}
\end{array}$$

A few comments are in order. We first note that Def. 7 is *compositional*; see e.g., [23] for reasons why this is desirable. It also assumes a bijective mapping between the denumerable sets LVAR and VARS; see synthesis for X , $\max X. \psi$ and $\min X. \psi$, where the logical variable X is converted to the process variable x . Although Def. 7 covers both SHML and CHML, the syntactic constraints of Def. 6 mean that synthesis for a formula ψ uses at most the first row and then either the first column (in the case of SHML) or the second column (in case of CHML). The conditional cases handle logically equivalent formulae, e.g., since $\llbracket \text{ff} \rrbracket = \llbracket \min X. \langle \alpha \rangle \text{ff} \rrbracket$ we have $\langle \text{ff} \rangle = \langle \min X. \langle \alpha \rangle \text{ff} \rangle = \text{no}$. In the case of conjunctions and disjunctions, these are essential to be able to generate sound monitors.

Example 8. Consider the CHML formula $\langle \alpha \rangle \text{tt} \vee \text{ff}$ (which is logically equivalent to $\langle \alpha \rangle \text{tt}$). A naive synthesis without the case checks would generate the monitor $\alpha. \text{yes} + \text{no}$ which is not only redundant, but *unsound* e.g., for $p = \alpha. \text{nil}$ we have both $\text{acc}(p, \alpha. \text{yes} + \text{no})$ and $\text{rej}(p, \alpha. \text{yes} + \text{no})$. Similar problems would manifest themselves for less obvious cases, e.g., $\langle \alpha \rangle \text{tt} \vee (\min X. \langle \alpha \rangle \text{ff}) \vee (\langle \alpha \rangle \min X. \text{ff})$. However, in all of these cases, our monitor synthesis of Def. 7 generates $\alpha. \text{yes}$. ■

Theorem 1 (Monitorability). $\varphi \in \text{MHML}$ implies φ is monitorable.

Proof. We show that for all $\varphi \in \text{MHML}$, $\text{mon}(\langle \varphi \rangle, \varphi)$ holds. □

Thm. 1 provides us with a simple syntactic check to determine whether a formula is monitorable; as shown earlier in Ex. 7, determining whether a formula is monitorable is in general non-trivial. Moreover, the proof of Thm. 1 (through Def. 7) provides us with an automatic monitor synthesis algorithm that is correct according to Def. 4.

Example 9. Since φ_4 from Ex. 2 is in MHML, we know it is monitorable. Moreover, we can generate the *correct* monitor $\langle \varphi_4 \rangle = \text{rec } x. (\text{req. ans. } x + \text{cls. no}) = m_2$ (from Ex. 3). Using similar reasoning, φ_7 (Ex. 4) is also monitorable, and a correct monitor for it is m_1 from Ex. 3. ■

7 Expressiveness

The results obtained in Sec. 6 beg the question of whether MHML is the *largest* monitorable subset of μHML . One way to provide an answer to this question would be to show that, in some sense, every monitor corresponds to a formula in MHML according to Def. 4. However, this approach quickly runs into problems since there are monitors, such as **yes** + **no**, that make little sense from the point of view of Def. 4. To this end, we prove a general (and perhaps surprising) result.

Theorem 2 (Multi-verdict Monitors and Monitoring). $\forall m \in \text{MON}$

$$(\exists t, u \in \text{ACT}^*. m \xrightarrow{t} \text{yes} \text{ and } m \xrightarrow{u} \text{no}) \text{ implies } \nexists \varphi \in \mu\text{HML}. \mathbf{mon}(m, \varphi).$$

Proof. By contradiction. Assume that $\exists \varphi. \mathbf{mon}(m, \varphi)$. Using t and u , we can construct the obvious process $p = t + u$, where $m \triangleleft p \xrightarrow{t} \text{yes} \triangleleft \text{nil}$ and $m \triangleleft p \xrightarrow{u} \text{no} \triangleleft \text{nil}$. We therefore have $\mathbf{acc}(p, m)$ and $\mathbf{rej}(p, m)$, and by monitor soundness (Def. 2), that $p \in \llbracket \varphi \rrbracket$ and $p \notin \llbracket \varphi \rrbracket$. This is clearly a contradiction. \square

Stated otherwise, Thm. 2 asserts that multi-verdict monitors are necessarily *unsound*, at least wrt. properties defined over processes (as opposed to logics defined over other domains such as traces, e.g., [11, 7, 14]). This also implies that, in order to answer the aforementioned question, it suffices to focus on *uni-verdict* monitors that flag either acceptances or rejections (but not both). In fact, a closer inspection of the synthesis algorithm of Def. 7 reveals that all the monitors generated using it are, in fact, uni-verdict.

We partition uni-verdict monitors into the obvious classes: *acceptance monitors*, AMON (using verdict **yes**), and *rejection monitors*, RMON (using **no**). In what follows, we focus our technical development on one monitor class, in the knowledge that the corresponding development for the other class is analogous.

Definition 8 (Rejection Expressive-Complete). A subset $\mathcal{L} \subseteq \mu\text{HML}$ is expressive-complete wrt. rejection monitors iff

$$\forall m \in \text{RMON}. \exists \varphi \in \mathcal{L} \text{ such that } \mathbf{mon}(m, \varphi).$$

We show that the language sHML (Def. 6) is rejection expressive-complete. We do so with the aid of a mapping function from a rejection monitor to a corresponding formula in sHML defined below. Def. 9 is fairly straightforward, thanks to the fact that we only need to contend with a single verdict. Again, it assumes a bijective mapping between the denumerable sets LVAR and VARS (as in the case of Def. 7). The mapping function is defined inductively on the structure of m where we note that (i) no translation is given for the monitor **yes** (since these are rejection monitors) and (ii) the base case **end** is mapped to formula **tt**, which contrasts with the mapping used in Def. 7.

Definition 9 (Rejection Monitors to sHML Formulae).

$$\begin{aligned} \langle\langle \text{no} \rangle\rangle &\stackrel{\text{def}}{=} \text{ff} & \langle\langle \text{end} \rangle\rangle &\stackrel{\text{def}}{=} \text{tt} & \langle\langle x \rangle\rangle &\stackrel{\text{def}}{=} X \\ \langle\langle \alpha.m \rangle\rangle &\stackrel{\text{def}}{=} [\alpha] \langle\langle m \rangle\rangle & \langle\langle m + n \rangle\rangle &\stackrel{\text{def}}{=} \langle\langle m \rangle\rangle \wedge \langle\langle n \rangle\rangle & \langle\langle \text{rec } x.m \rangle\rangle &\stackrel{\text{def}}{=} \max X. \langle\langle m \rangle\rangle \end{aligned}$$

Proposition 2. *sHML is Rejection Expressive-Complete.*

Proof. We show that for all $m \in \text{RMON}$, $\mathbf{mon}(m, \langle\langle m \rangle\rangle)$ holds. \square

Definition 10 (Acceptance Expressive-Complete). *Language $\mathcal{L} \subseteq \mu\text{HML}$ is expressive-complete wrt. acceptance monitors iff $\forall m \in \text{AMON}. \exists \varphi \in \mathcal{L}$ such that $\mathbf{mon}(m, \varphi)$.*

Proposition 3. *cHML is Acceptance Expressive-Complete.*

Equipped with Prop. 2 and Prop. 3, it follows that μHML is expressive complete wrt. uni-verdict monitors.

Definition 11 (Expressive-Complete). *$\mathcal{L} \subseteq \mu\text{HML}$ is expressive-complete wrt. uni-verdict monitors, iff $\forall m \in \text{AMON} \cup \text{RMON}. \exists \varphi \in \mathcal{L}. \mathbf{mon}(m, \varphi)$.*

Theorem 3. *μHML is Expressive-Complete.*

Proof. Follows from Prop. 2 and Prop. 3 \square

We are now in a position to prove the result alluded to at the beginning of this section, namely that μHML is the largest monitorable subset of μHML up to logical equivalence, i.e., Thm. 4. First, however, we define what we understand by language inclusion up to formula semantic equivalence, Def. 12.

Definition 12 (Language Inclusion). *For all $\mathcal{L}_1, \mathcal{L}_2 \in \mu\text{HML}$*

$$\mathcal{L}_1 \sqsubseteq \mathcal{L}_2 \stackrel{\text{def}}{=} \forall \varphi_1 \in \mathcal{L}_1. \exists \varphi_2 \in \mathcal{L}_2 \text{ such that } \llbracket \varphi_1 \rrbracket = \llbracket \varphi_2 \rrbracket$$

We also prove the following important proposition, that gives an upper bound to the expressiveness of languages satisfying monitorability properties.

Proposition 4. *For any $\mathcal{L} \subseteq \mu\text{HML}$:*

1. *($\forall \varphi \in \mathcal{L}. \exists m \in \text{AMON}. \forall p. (\mathbf{acc}(p, m) \text{ iff } p \in \llbracket \varphi \rrbracket)$) implies $\mathcal{L} \sqsubseteq \text{cHML}$.*
2. *($\forall \varphi \in \mathcal{L}. \exists m \in \text{RMON}. \forall p. (\mathbf{rej}(p, m) \text{ iff } p \notin \llbracket \varphi \rrbracket)$) implies $\mathcal{L} \sqsubseteq \text{sHML}$.*

Proof. We prove the first clause; the second clause is analogous. Assume $\varphi \in \mathcal{L}$. We need to show that $\exists \pi \in \text{cHML}$ such that $\llbracket \varphi \rrbracket = \llbracket \pi \rrbracket$. For φ we know that

$$\exists m \in \text{AMON}. (\forall p. (\mathbf{acc}(p, m) \text{ iff } p \in \llbracket \varphi \rrbracket)). \quad (1)$$

By Prop. 3, for the monitor m used in (1), we also know

$$\exists \pi \in \text{cHML}. (\forall p. (\mathbf{acc}(p, m) \text{ iff } p \in \llbracket \pi \rrbracket)). \quad (2)$$

Assume an arbitrary $p \in \llbracket \varphi \rrbracket$. By (1) we obtain $\mathbf{acc}(p, m)$, and by (2) we obtain $p \in \llbracket \pi \rrbracket$. Thus $\llbracket \varphi \rrbracket \subseteq \llbracket \pi \rrbracket$. Dually, we can also reason that $\llbracket \pi \rrbracket \subseteq \llbracket \varphi \rrbracket$. \square

Theorem 4 (Completeness). *($\mathcal{L} \subseteq \mu\text{HML}$ is monitorable) implies $\mathcal{L} \sqsubseteq \mu\text{HML}$.*

Proof. Since \mathcal{L} is monitorable, by Def. 5 and Def. 4 we know:

$$\forall \varphi \in \mathcal{L}. \exists m \text{ such that } \mathbf{smon}(m, \varphi) \text{ and } \mathbf{cmon}(m, \varphi) \quad (3)$$

By Thm. 2, we know that every m used in (3) is uni-verdict. This means that we can partition the formulae in \mathcal{L} into two disjoint sets $\mathcal{L}_{\text{acc}} \uplus \mathcal{L}_{\text{rej}}$ where:

$$(\forall \varphi \in \mathcal{L}_{\text{acc}}. \exists m \in \text{AMON}. \forall p. (\mathbf{acc}(p, m) \text{ iff } p \in \llbracket \varphi \rrbracket)) \quad (4)$$

$$(\forall \varphi \in \mathcal{L}_{\text{rej}}. \exists m \in \text{RMON}. \forall p. (\mathbf{rej}(p, m) \text{ iff } p \notin \llbracket \varphi \rrbracket)) \quad (5)$$

By (4), (5) and Prop. 4 we obtain $\mathcal{L}_{\text{acc}} \sqsubseteq \text{CHML}$ and $\mathcal{L}_{\text{rej}} \sqsubseteq \text{SHML}$ *resp.*, from which the required result follows. \square

Thm. 2 and Thm. 4 constitute powerful results *wrt.* the monitorability of our branching-time logic. Completeness, Thm. 4, guarantees that limiting oneself to the syntactic subset MHML does not hinder the *expressive power* of the specifier when formulating monitorable properties. Alternatively, one could also determine whether a formula is monitorable by rewriting it as a logically equivalent formula in MHML.⁵ This would enable a verification framework to decide whether to check for a μHML property at runtime, or resort to more expressive (but expensive means) otherwise. Whenever the property is monitorable, Thm. 2 guarantees that a uni-verdict monitor is the best monitor that we can synthesise. This is important since multi-verdict monitor constructions, such as those in [7], generally carry *higher overheads* than uni-verdict monitors.

Example 10. By virtue of Thm. 4, we can conclude that properties $\varphi_1, \varphi_2, \varphi_3, \varphi_5$ and φ_6 from Ex. 2 are all non-monitorable properties according to Def. 5, since no logically equivalent formulae in MHML exist. Arguably, the problem of establishing logical equivalence through syntactic manipulation of formulae is easier to determine and automate, when compared to direct reasoning about the semantic definitions of monitorability and those of the *resp.* properties; recall that Def. 5 (Monitorability) — through Def. 2 and Def. 3 — universally quantifies over all processes, which generally poses problems for automation.

For instance, in the case of φ_1 , we could use Thm. 2 to substantially reduce the search space of our witness monitor to the uni-verdict ones, but this still leaves us with a lot of work to do. Specifically, we can reason that the witness *cannot* be an *acceptance monitor*, since it would need to accept process nil , which implies that it must erroneously also accept the process cls.nil (using reasoning similar to that used in Ex. 6). It is less straightforward to argue that the witness *cannot* be a *rejection monitor* either. We argue towards a contradiction by assuming that such a monitor exists. Since it is violation-complete (Def. 3) it should reject the process $\text{req.nil} + \text{cls.nil}$ since this process does not satisfy φ_1 : by Prop. 1 we know that it can do so along either of the traces ϵ, req or cls . If it rejects it along ϵ , then it also rejects the satisfying process nil ; if it rejects along

⁵ The problem of determining whether a (general) formula is logically equivalent to one in MHML is decidable in exponential time — probably EXPTIME complete.

trace `req`, it also rejects the satisfying process `req.ans.nil`; finally, if it rejects it along `cls`, it must also reject the satisfying process `req.ans.nil + cls.nil`. Thus, the monitor must be unsound, meaning that it cannot be a rejection monitor. ■

8 Conclusion

We have investigated monitorability aspects of a branching-time logic called μ HML, which impinges on what properties can be verified at runtime. It extends and generalises prior work carried out in the context of reactive systems modelled as LTSs [16]. The concrete contributions of the paper are:

1. An operational definition of monitorability, Def. 4, specified over an instrumentation relation unifying the individual behaviour of processes and monitors, Fig. 4, which is used to define monitorable subsets of μ HML, Def. 5.
2. The identification of a subset of μ HML, Def. 6, that is shown to be monitorable, Thm. 1, and also maximally expressive, Thm. 4, wrt. Def. 5.
3. A result asserting that, wrt. Def. 4, uni-verdict monitors suffice for monitoring branching-time properties, Thm. 2.

Future Work: It is worth exploring other definitions of monitorability apart from that of Def. 4, and determining how this affects the monitorable subset of μ HML identified in this work. For instance, one could relax the conditions of Def. 4 by only requiring soundness (Def. 2), or require more stringent conditions wrt. verdicts and monitor non-determinism; see [16] for a practical motivation of this. Moreover, monitorability is also largely dependent on the underlying instrumentation relation used; there may be other sensible relations apart from the one defined in Fig. 4 that are worth investigating within this setting.

A separate line of research could investigate manipulation techniques that decompose formulae into monitorable components. For instance, reformulating a generic formula φ as the disjunction $\phi \vee \pi$ (recall $\pi \in \text{cHML}$) could allow for a *hybrid* verification approach that distributes the load between the pre-deployment phase and the runtime phase whereby we model-check for the satisfaction of a system wrt. ϕ and, if this fails, runtime verify the system wrt. π .

Related Work: In [23], monitorability is defined for *formulae defined over traces* (e.g., LTL) whenever the formula semantics does not contain *ugly* prefixes; an ugly prefix is a trace from which *no* finite extension will ever lead to a conclusive verdict. Falcone *et al.* [14] revisit this classical definition, extending it to the Safety-Progress property classification, while proposing an alternative definition in terms of the structure of the recognising Streett Automata of the *resp.* property. Although our definition is cast within a different setting (a logic over processes), and has a distinct operational flavour in terms of monitored system executions, it is certainly worthwhile to try to reconcile the different definitions.

The logic μ HML has been previously studied from a linear-time perspective in [2, 9], in order to find subsets that characterise may/must testing equivalences. Although tests are substantially different from our monitor instrumentations, the logic subsets identified in [2, 9] are related to (albeit different from) mHML.

References

1. J. R. Abrial. *Modeling in Event-B: System and Software Engineering*. Cambridge University Press, 2010.
2. L. Aceto and A. Ingólfssdóttir. Testing Hennessy-Milner Logic with Recursion. In *FoSSaCS'99*, pages 41–55. Springer, 1999.
3. L. Aceto, A. Ingólfssdóttir, K. G. Larsen, and J. Srba. *Reactive Systems: Modelling, Specification and Verification*. Cambridge Univ. Press, New York, NY, USA, 2007.
4. H. Barringer, A. Goldberg, K. Havelund, and K. Sen. Rule-based runtime verification. In *VMCAI*, volume 2937 of *LNCS*, pages 44–57. Springer, 2004.
5. A. Bauer, M. Leucker, and C. Schallhart. The good, the bad, and the ugly, but how ugly is ugly? In *RV*, volume 4839 of *LNCS*, pages 126–138. Springer, 2007.
6. A. Bauer, M. Leucker, and C. Schallhart. Comparing LTL semantics for runtime verification. *Logic and Comput.*, 20(3):651–674, 2010.
7. A. Bauer, M. Leucker, and C. Schallhart. Runtime verification for LTL and TLTL. *TOSEM*, 20(4):14, 2011.
8. I. Cassar and A. Francalanza. On Synchronous and Asynchronous Monitor Instrumentation for Actor Systems. In *FOCLASA*, volume 175, pages 54–68, 2014.
9. A. Cerone and M. Hennessy. Process behaviour: Formulae vs. tests. In *EXPRESS*, volume 41 of *EPTCS*, pages 31–45, 2010.
10. E. Chang, Z. Manna, and A. Pnueli. Characterization of temporal property classes. In *ALP LNCS*, pages 474–486. Springer-Verlag, 1992.
11. C. Cini and A. Francalanza. An LTL Proof System for Runtime Verification. In *TACAS*, volume 9035, pages 581–595. Springer, 2015.
12. C. Colombo, G. Pace, and G. Schneider. LARVA — Safer monitoring of Real-Time Java programs (Tool paper). In *SEFM*, pages 33–37, 2009.
13. C. Eisner, D. Fisman, J. Havlicek, Y. Lustig, A. McIsaac, and D. V. Campenhout. Reasoning with temporal logic on truncated paths. In *CAV*, volume 2725 of *LNCS*, pages 27–39. Springer, 2003.
14. Y. Falcone, J.-C. Fernandez, and L. Mounier. What can you verify and enforce at runtime? *STTT*, 14(3):349–382, 2012.
15. A. Francalanza, A. Gauci, and G. J. Pace. Distributed System Contract Monitoring. *JLAP*, 82(5-7):186–215, 2013.
16. A. Francalanza and A. Seychell. Synthesising Correct concurrent Runtime Monitors. *Formal Methods in System Design (FMSD)*, pages 1–36, 2014.
17. M. Geilen. On the Construction of Monitors for Temporal Logic Properties. In *RV*, volume 55 of *ENTCS*, pages 181–199, 2001.
18. C. A. R. Hoare. *Communicating sequential processes*. Prentice-Hall, 1985.
19. D. Kozen. Results on the propositional μ -calculus. *TCS*, 27:333–354, 1983.
20. M. Leucker and C. Schallhart. A brief account of Runtime Verification. *JLAP*, 78(5):293 – 303, 2009.
21. Z. Manna and A. Pnueli. Completing the Temporal Picture. *TCS*, 83(1):97–130, 1991.
22. R. Milner. *A Calculus of Communicating Systems*. Springer, 1982.
23. A. Pnueli and A. Zaks. Psl model checking and run-time verification via testers. In *FM*, pages 573–586. Springer, 2006.
24. K. Sen, G. Rosu, and G. Agha. Generating optimal linear temporal logic monitors by coinduction. In *ASIAN*, LNCS, pages 260–275. Springer, 2004.
25. detectEr Project. <http://www.cs.um.edu.mt/svrg/Tools/detectEr/>.