

Monitorability for the Hennessy-Milner Logic with Recursion

Adrian Francalanza · Luca Aceto · Anna Ingolfsdottir

Abstract We study μ HML, a branching-time logic with least and greatest fixpoints, from a runtime verification perspective. The logic may be used to specify properties of programs whose behaviour may be expressed as a labelled transition system. We establish which subset of this logic can be monitored for at runtime by merely observing the runtime execution of a program. A monitor-synthesis algorithm is defined for this subset, where it is shown that the resulting synthesised monitors correctly perform the required analysis from the observed behaviour. We also prove completeness results with respect to this logical subset that show that, up to logical equivalence, no other properties apart from those identified can be monitored for and verified at runtime.

The research of L. Aceto and A. Ingolfsdottir was partly supported by the project 001-ABEL-CM-2013 of the NILS Science and Sustainability Programme. The research of all three authors was also supported by the project Theoretical Foundations of Monitorability (nr. 163406-051) of the Icelandic Research Fund.

Adrian Francalanza
Dept. of Computer Science,
ICT, University of Malta, Msida, Malta.
Tel.: +356-2340-2745
Fax: +356-2132-0539
E-mail: adrian.francalanza@um.edu.mt

Luca Aceto
School of Computer Science,
Reykjavik University, Reykjavik, Iceland
Tel. : +354 599 6419
Fax. : +354 599 6301
E-mail: luca@ru.is

Anna Ingolfsdottir
School of Computer Science,
Reykjavik University, Reykjavik, Iceland
Tel. : +354 599 6319
Fax. : +354 599 6301
E-mail: anna@ru.is

1 Introduction

Runtime Verification (RV) [31,22] is a lightweight verification technique whereby the *execution of a system* is analysed with the aim of inferring correctness with respect to some property. The technique is often used to mitigate scalability issues such as state-explosion problems, typically associated with more traditional verification techniques like Model Checking. RV is also useful in settings where the full execution state space of a system is not accessible or available prior deployment, and the only means of analysing the behaviour of a system is by observing its current system run. Despite its advantages, the technique is generally limited when compared to other verification techniques that consider system behaviour beyond its current execution [32,17]. For instance, *online* RV incrementally analyses *partial* executions (up to the current execution point of the system) which limits its applicability to satisfaction analyses relating to correctness properties that describe complete (*i.e.*, potentially infinite) executions.

There are broadly two approaches in the field of RV that are used 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) [8,25], or else *redefines* the semantics of existing logics (*e.g.*, LTL) so as to reflect the limitations of the runtime setting [21,10,12,11]. The second approach is to leave the semantics of the specification logic *unchanged*, and then identifying which *subsets* of the logic can be verified at runtime, with respect to a specific monitoring setup [27,35,18,26,20,38].

Both approaches have their merits. For instance, the first approach is generally more popular and tends to produce specifications that are closely related to the monitors that check for them (*e.g.*, RE and automata in [34,8,25]), thus facilitating aspects such as the study of monitor correctness (which is well-understood in certain settings). 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 to synthesise a monitor that checks 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 combined approach, where some parts of a property are verified using RV and other parts of the property are checked using other techniques, along the lines of [4].

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. Prosaically, when specifying the correct behaviour expected of a system (typically during the design stage of development) one should ideally be unencumbered by how this correctness specification is checked for in the target system. Instead, the focus at this development stage should be on the expressiveness of the specification formalism used. Indeed, there are many cases where the verification strategy to be adopted becomes apparent later on in the development process or else is changed due to evolving requirements of the software project (*e.g.*, the development team may decide to use a proprietary software component whose source code is unavailable). Alternatively, the project requirements may change in such a way so as to require more expressive logics to be able to specify the augmented correctness requirements.

This work adheres to the second approach discussed above. In particular, it revisits the Hennessy-Milner Logic with recursion [30,3], μ HML, a reformulation of the modal μ -calculus [29], used to describe correctness properties of reactive system. We choose to conduct our study in this logic for a number of reasons. The logic has been shown to be a very powerful and expressive formalism for describing behavioural properties of systems; it is more expressive than widely used temporal logics such as CTL and CTL* [19,7] which suggests that the findings of this work can be carried over to these logics. In addition, our work can serve to complete the picture for existing work in the context of RV: fragments of μ HML have already been adapted for detectEr [37], an RV tool used to specify and verify at runtime actor-based reactive systems [26,13–15], and constructs from the modal μ -calculus have been used in other RV tools such as Eagle [9]. In this study, we consider the logic in its entirety, and investigate the monitorability of fragments of the logic with respect to an operational definition of a general class of monitors that employ both acceptance and rejection verdicts [12,22]. It turns out that our results extend the class of monitorable μ HML properties used in [26] and establish monitorability upper bounds for this logic. We also present new results that relate the utility of multi-verdict monitors for logics that are defined over program execution graphs, in contrast to other work on monitorability that considers logics and properties defined over traces [12,22,18]. The concrete contributions of the paper are:

1. An operational definition of what it means for a formula in a branching-time logic to be monitored by a specific monitor, Definition 4. Underpinning this definition is an instrumentation relation unifying the individual behaviour of processes and monitors as a single monitored system entity, given in Figure 4 that complements other foundational work on monitor behavior [23,14].
2. A definition that formalises when a logical formula is considered to be monitorable in this setting, Definition 5, that is then lifted to logical languages, used to define monitorable subsets of μ HML.
3. The identification of a subset of μ HML, Definition 6, that is shown to be monitorable according to Definition 5 (see Theorem 1) but also maximally expressive with respect to this monitorability definition, Theorem 4. This means that there are no other (semantically equivalent) properties in μ HML that can be monitored for in the setting that we consider.
4. A result asserting that with respect to the monitor setting of Definition 4, uni-verdict monitors, *i.e.*, monitors with just an acceptance or a rejection verdict, suffice to attain the full expressive monitoring power, Theorem 2. We assert that this property is inherent to branching-time specifications in general.

To the best of our knowledge, this is one of the first bodies of work (along with [38]) investigating the limits of RV with respect to the *branching-time* dimension of logical specifications (*i.e.*, properties that talk about the *execution graph* of a program). Other work pertaining to the aforementioned second approach tend to follow a more traditional approach, focussing on *linear-time* logics defined over *execution traces*, exploring RV's limits along the linear-time dimension, see *e.g.*, [18].

This article is a full version of the extended abstract published as the conference paper [24]; it includes the detailed definitions, expanded explanations, more detailed examples, and complete proofs. We have also updated the related work with recent literature. The rest of the paper is structured as follows. Section 2

Syntax

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

Dynamics

$$\begin{array}{c}
\text{ACT} \frac{}{\alpha.p \xrightarrow{\alpha} p} \quad \text{SELL} \frac{p \xrightarrow{\mu} p'}{p + q \xrightarrow{\mu} p'} \quad \text{SELR} \frac{q \xrightarrow{\mu} q'}{p + q \xrightarrow{\mu} q'} \\
\text{SYN} \frac{p \xrightarrow{\alpha} p' \quad q \xrightarrow{\bar{\alpha}} q'}{p \parallel q \xrightarrow{\tau} p' \parallel q'} \quad \text{PARL} \frac{p \xrightarrow{\mu} p'}{p \parallel q \xrightarrow{\mu} p' \parallel q} \quad \text{PARR} \frac{q \xrightarrow{\mu} q'}{p \parallel q \xrightarrow{\mu} p \parallel q'} \\
\text{REC} \frac{}{\text{rec } x.p \xrightarrow{\tau} p[\text{rec } x.p/x]} \quad \text{SCP} \frac{p \xrightarrow{\mu} p' \quad \mu \notin \{\alpha, \bar{\alpha}\}}{\text{new } \alpha.p \xrightarrow{\mu} \text{new } \alpha.p'}
\end{array}$$

Fig. 1 A Model for describing Systems

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

2 The Model

We describe systems abstractly as Labelled Transition Systems (LTSs) [3, 29]. An LTS is a triple $\langle \text{PROC}, (\text{ACT} \cup \{\tau\}), \longrightarrow \rangle$ consisting of a set of states, $p, q, \dots \in \text{PROC}$, a set of actions, $\alpha, \bar{\alpha}, \beta, \bar{\beta}, \dots \in \text{ACT}$ with distinguished silent action τ , and a transition relation, $\longrightarrow \subseteq (\text{PROC} \times (\text{ACT} \cup \{\tau\}) \times \text{PROC})$. We assume that the set of actions ACT comes equipped with an involution function $\bar{\alpha}$, where $\bar{\bar{\alpha}} = \alpha$, and where $\bar{\alpha}$ is considered to be the complement of α . We also assume that $\mu \in \text{ACT} \cup \{\tau\}$, that $\tau \notin \text{ACT}$ and refer to the actions in ACT as *visible* actions (as opposed to the silent action τ).

LTS states can be expressed as processes, PROC , from a standard variant of CCS [33] as defined by the syntax in Figure 1. Assuming a specific set of (visible) actions ACT and a denumerable set of (recursion) variables $x, y, z \in \text{VARS}$, processes are defined as either the inactive process nil , a prefixed process by an action α , a mutually-exclusive choice between two processes, a parallel composition of two processes, a scoping of an action in a process, or a recursive process. The process construct $\text{rec } x.p$ acts as a *binder* for x in p whereas the process construct $\text{new } \alpha.p$ acts as a binder for actions $\alpha, \bar{\alpha}$ in p ; we work up to alpha-conversion of bound variables and scoped actions. All recursive processes are assumed to be guarded,




Fig. 2 A depiction of the system in Example 1

meaning that all occurrences of bound variables occur under an action prefix (either directly or indirectly). Closed terms are processes where all occurrences of variables are bound. Occasionally, we elide inactive processes when describing systems and write p instead of $\text{nil} \parallel p$ or $p \parallel \text{nil}$ and $\alpha.\beta.\text{nil}$ as $\alpha.\beta$; in order to lighten our presentation we also implicitly work up to associativity and commutativity of parallel compositions where, for example, $p \parallel q$ is some times treated as $q \parallel p$.¹

The dynamic behaviour of a process is described by the transition rules of Figure 1, defined over the closed and guarded terms in PROC. The suggestive notation $p \xrightarrow{\mu} p'$ denotes $(p, \mu, p') \in \longrightarrow$; we also write $p \not\xrightarrow{\mu}$ to denote $\neg(\exists p'. p \xrightarrow{\mu} p')$. The rules in Figure 1 are standard. For example, $p_1 + p_2 \xrightarrow{\mu} q$ if either $p_1 \xrightarrow{\mu} q$ or $p_2 \xrightarrow{\mu} q$ by rules SELL or SELR; dually, $p_1 \parallel p_2 \xrightarrow{\mu} q$ if $p_1 \xrightarrow{\mu} p'_1$ and $q = p'_1 \parallel p_2$ for some p'_1 , or $p_2 \xrightarrow{\mu} p'_2$ and $q = p_1 \parallel p'_2$ for some p'_2 . We note that rule SYN describes the synchronisation of two parallel processes on complementing actions whereas rule SCP restricts visible actions to those that are not scoped by α in $\text{new } \alpha.p$. We employ the usual notation for weak transitions and write $p \Longrightarrow p'$ in lieu of $p(\xrightarrow{\tau})^* p'$ 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 sequences of transitions $p \xRightarrow{\alpha_1} \dots \xRightarrow{\alpha_n} p_n$ as $p \xRightarrow{t} p_n$, where $t = \alpha_1, \dots, \alpha_n$. For more details the reader is invited to consult standard texts such as [33, 3].

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 .

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

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.

$$\begin{aligned} q &= \text{rec } x.(\text{req}.\overline{\text{ans}}.x + \text{cls}.\text{nil} + (\text{rec } y.\text{req}.\overline{\text{ans}}.y)) \\ r &= \text{req}.\overline{\text{ans}}.\text{rec } x.(\text{req}.\overline{\text{ans}}.x + \text{cls}.\text{nil}) \end{aligned}$$

¹ These are normally expressed as structural equivalence rules such as $p \parallel \text{nil} \equiv p$ and $p \parallel q \equiv q \parallel p$ in standard CCS. We elide them here to alleviate our exposition.

We will use the processes p, q and r to exemplify some of the technical notions to follow. Abstractly, the respective LTSs denoted by these processes may be represented by the graphs in Figure 2, where the nodes correspond to processes and the arcs correspond to weak transitions, $\xrightarrow{\alpha}$. Note that these graphs do not describe the full branching structure of the respectively processes and abstract away from τ -derivatives such as $(\text{req}.\overline{\text{ans}}.p + \text{cls}.\text{nil})$ (in the case of p). This omission is intensional and the intuition conveyed by these graphs, namely finite depictions of *sets of transitions* $p' \xrightarrow{t} p''$, suffices for the intuitions underpinning the logical specifications that follow. ■

Example 2 Consider the system described by the process p' below:

$$p' = \text{rec } x. \left((\text{req}.\text{work}.\text{nil} + \text{cls}.\text{nil}) \parallel \overline{\text{work}}.\overline{\text{ans}}.x \right)$$

By abstracting away from τ -transitions, we can also use the finite graph for p from Figure 2 to describe the behaviour of the process $\text{new work}.p'$, where the action work is locally scoped. In particular, such a graph abstracts away from the τ -transitions (1) and (2) below.

$$\text{new work}.p' \xrightarrow{\tau} \text{new work}.\left((\text{req}.\text{work}.\text{nil} + \text{cls}.\text{nil}) \parallel \overline{\text{work}}.\overline{\text{ans}}.p'\right) \quad (1)$$

$$\begin{aligned} &\xrightarrow{\text{req}} \text{new work}.\left(\text{work}.\text{nil} \parallel \overline{\text{work}}.\overline{\text{ans}}.p'\right) \\ &\xrightarrow{\tau} \text{new work}.\left(\overline{\text{ans}}.p'\right) \quad (2) \\ &\xrightarrow{\overline{\text{ans}}} \text{new work}.p' \end{aligned}$$

Importantly, *not* all processes describe a finite execution graph. As a simple demonstration of this, consider the process variants p'' and p''' below:

$$\begin{aligned} p'' &= \text{rec } x. \text{new work}.\left(\overline{\text{work}}.\text{nil} \parallel (\text{req}.\overline{\text{ans}}.x + \text{cls}.\text{nil})\right) \\ p''' &= \text{rec } x. \left(\text{req}.\overline{\text{ans}}.(x \parallel x) + \text{cls}.\text{nil}\right) \end{aligned}$$

In the case of p'' , we can easily observe the sequence of non-repeating states where new parallel copies of $\overline{\text{work}}.\text{nil}$ are spawned (we rename the bound name work to work' to highlight distinct bindings):

$$\begin{aligned} p'' &\xrightarrow{\text{req}.\overline{\text{ans}}} \text{new work}.\left(\overline{\text{work}}.\text{nil} \parallel p''\right) \\ &\xrightarrow{\text{req}.\overline{\text{ans}}} \text{new work}.\left(\overline{\text{work}}.\text{nil} \parallel \text{new work}'.\left(\overline{\text{work}}'.\text{nil} \parallel p''\right)\right) \xrightarrow{\text{req}.\overline{\text{ans}}} \dots \end{aligned}$$

We can also observe a similar expansion in the case of p''' . For instance, consider the execution trace below:

$$p''' \xrightarrow{\text{req}.\overline{\text{ans}}} p''' \parallel p''' \xrightarrow{\text{req}.\overline{\text{ans}}} p''' \parallel p''' \parallel p''' \xrightarrow{\text{req}.\overline{\text{ans}}} \dots$$

In fact, the execution graph of p''' is far more complicated than that since req and $\overline{\text{ans}}$ from the generated parallel copies of p''' can arbitrarily be interleaved. ■

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

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

Fig. 3 μHML Syntax and Semantics**3 The Logic**

The logic μHML [30,3] 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 Figure 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 $\max X. \varphi$ respectively bind free instances of the logical variable X in φ , inducing the usual notions of open/closed formulae and formula equality up to alpha-conversion.

Formulae are interpreted over the process powerset domain, $S \in \mathcal{P}(\text{PROC})$. The semantic definition of Figure 3 is given for *both* open and closed formulae and employs an environment from logical variables to sets of processes, $\rho \in \text{LVAR} \rightarrow \mathcal{P}(\text{PROC})$. This permits an inductive definition of $\llbracket \varphi, \rho \rrbracket$, the set of processes satisfying the formula φ with respect to an environment ρ , based on the structure of the formula. For instance, in Figure 3, the semantic meaning of a variable X with respect to 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 with respect to set-inclusion, \subseteq , which guarantees the existence of least and largest solutions for the recursive formulae of the logic; as usual, these can be respectively 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 closed formula φ whenever $p \in \llbracket \varphi \rrbracket$, and *violates* a formula whenever $p \notin \llbracket \varphi \rrbracket$.

Example 3 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 . Consider the formulae defined as follows:

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

Formula φ_1 denotes a *liveness* property describing processes that stop offering the action cls after some number of serviced request, *i.e.*, after performing some sequence of actions in $(\text{req}.\overline{\text{ans}})^*$ —processes q and r from Example 1 satisfy this property but p does not. Concretely, in the case of q we can observe the computation $q \xrightarrow{\text{req}.\overline{\text{ans}}} q_4$ and $q_4 \in \llbracket [\text{cls}] \text{ff} \rrbracket$ since $q_4 \not\xrightarrow{\text{cls}}$, whereas in the case of r we immediately have $r \xrightarrow{\text{cls}}$. Intuitively, p does not satisfy φ_1 because we have $p \xrightarrow{\text{cls}}$ and whenever $p \xrightarrow{\text{req}.\overline{\text{ans}}} p'$ for some p' , then it necessarily the case that $p' = p$ again (or its unfolding $\text{req}.\overline{\text{ans}}.p + \text{cls}.\text{nil}$). If we change the fixpoint of φ_1 into a maximal one, *i.e.*, φ_2 , the resulting formula would include p in the property as well. This is because $\llbracket \varphi_2 \rrbracket$ also includes processes that exhibit the *infinite* trace $(\text{req}.\overline{\text{ans}})^*$, which process p can display by following the looping transition sequence $p \xrightarrow{\text{req}.\overline{\text{ans}}} p$.

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 requests, *i.e.*, after having performed any trace in $(\text{req}.\overline{\text{ans}})^*$ (the maximal fixpoint allows these processes to perform an infinite number of serviced requests). Process p satisfies property φ_3 , due to an analogous argument to the one above justifying $p \notin \llbracket \varphi_1 \rrbracket$. However q and r do not satisfy this property because not all execution paths satisfy this property. In the case of process q , although paths along the looping sequence $q \xrightarrow{\text{req}.\overline{\text{ans}}} q$ *validate* this property (since q itself can perform the (weak) transition $q \xrightarrow{\text{cls}}$), it *violates* this property along a different path, such as in the case of $q \xrightarrow{\text{req}.\overline{\text{ans}}} q_4$. Formula φ_4 is analogous but *prohibits* the cls action after any number of serviced requests; we revisit this formula again in Example 7.

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

Formula φ_5 is satisfied by processes that, after one serviced request, $\text{req}.\overline{\text{ans}}$, exhibit a maximal² transition sequence in $(\text{req}.\overline{\text{ans}})^* \cup (\text{req}.\overline{\text{ans}})^\omega$ that never offers action cls . Process q satisfies this property through the path $q \xrightarrow{\text{req}.\overline{\text{ans}}} q_4$, whereas processes p and r do not satisfy such a property. Formula φ_6 describes processes that along *all* serviced request sequences, *i.e.*, traces in $(\text{req}.\overline{\text{ans}})^*$, eventually reach a state where they offer action cls . All the processes in Figure 2 satisfy this property: p and q satisfy the criteria immediately, whereas process r satisfies it for

² A transition sequence is maximal 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

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

Fig. 4 Monitors and Instrumentation

$(\text{req}.\overline{\text{ans}})^n$ sequences where $n \geq 1$. At an intuitive level, property φ_6 is, in some sense, analogous to the LTL formula $\langle \text{req} \rangle \langle \overline{\text{ans}} \rangle \text{tt } \mathbf{U} \langle \text{cls} \rangle \text{tt}$. ■

4 Monitors and instrumentation

Monitors are often expressed in terms of automata [8, 22] and may therefore also be viewed as LTSs, through the syntax of Figure 4. The syntax is similar to a sub-language of that for processes given in Figure 1, with the exception that the inert process nil is replaced by three *verdict* constructs, yes, no and end, respectively denoting acceptance, rejection and termination (*i.e.*, an inconclusive outcome). Monitor behaviour is similar to that of processes for the common constructs; see the transition rules in Figure 4. The only new transition rule concerns verdicts, MVER, stating that a verdict may transition with *any* $\alpha \in \text{ACT}$ and go back to the same state, thereby modelling the requirement that verdicts are *irrevocable*.

Figure 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 operations over LTSs [33, 28]. Specifically, rule IMON states that if a process can transition with action α and the respectively 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 (immediately) follow such a transition, $m \not\xrightarrow{\alpha}$, and (since it is already stable) there is no chance of it being able to do so after any sequence internal actions, $m \xrightarrow{\tau}$, the instrumentation forces it to terminate with an in-

conclusive verdict, `end`, while the process is allowed to proceed unaffected; see rule `ITER`. We remark that since `end` can produce *any* action, future transitions by p are still allowed while the terminated monitor maintains its state, using the rule `IMON`. Rules `IASYP` and `IASYM` allow monitors and processes to transition independently with respect to internal moves. This means that when a monitor cannot match a process action, but can transition silently, the instrumentation allows it to do so, and the matching check is then postponed to the τ -derivative monitor.

Example 4 Consider the monitor m defined below, which reaches an *acceptance* verdict after observing the action `req` followed by the action `ans`, but reaches a rejection verdict when it observes the action `ans` first.

$$m = (\text{req}.\overline{\text{ans}}.\text{yes}) + (\overline{\text{ans}}.\text{no})$$

When it is instrumented with process p from Example 1, we observe the following behaviour for the monitored system whereby on line (**) the monitor preserves the `yes` verdict for all remaining transitions.

$$\begin{aligned} m \triangleleft p &\xrightarrow{\tau} m \triangleleft \text{req}.\overline{\text{ans}}.p + \text{cls}.\text{nil} && \text{using IASYP} \\ &\xrightarrow{\text{req}} \overline{\text{ans}}.\text{yes} \triangleleft \overline{\text{ans}}.p && \text{using IMON} \\ &\xrightarrow{\overline{\text{ans}}} \text{yes} \triangleleft p && \text{using IMON} \\ &\xrightarrow{\text{req}.\overline{\text{ans}}} \text{yes} \triangleleft p \xrightarrow{\text{req}.\overline{\text{ans}}} \text{yes} \triangleleft p \xrightarrow{\text{req}.\overline{\text{ans}}} \dots && (**) \end{aligned}$$

When it is instrumented with a different system, such as $\overline{\text{ans}}.p$, the monitor yields a different (persistent) verdict,

$$m \triangleleft \overline{\text{ans}}.p \xrightarrow{\overline{\text{ans}}} \text{no} \triangleleft p \xrightarrow{\text{req}.\overline{\text{ans}}} \text{no} \triangleleft p \xrightarrow{\text{req}.\overline{\text{ans}}} \dots$$

Importantly, when m is instrumented with the process `req.ans.nil` we observe the following monitored system behaviour below. In the second transition, rule `ITER` is used because the emitted action, `ans` does not correspond with any of the actions the monitor is expecting at that stage, namely `ans` (note the bar on one of the latter action, `ans`, which makes it *distinct* from the former action `ans` from a monitoring perspective).

$$\begin{aligned} m \triangleleft \text{req.ans.nil} &\xrightarrow{\text{req}} (\overline{\text{ans}}.\text{yes}) \triangleleft (\text{ans}.nil) && \text{using IMON} \\ &\xrightarrow{\text{ans}} \text{end} \triangleleft \text{nil} && \text{using ITER} \end{aligned}$$

It is important that a monitor is terminated (with an inconclusive verdict) whenever it cannot match the exhibited trace event by the process, since leaving the monitor state unaltered may lead to unwanted detections. Consider the instrumentation of monitor m with the process `req.cls.ans.nil` and let us analyse the runtime behaviour below:

$$\begin{aligned} m \triangleleft \text{req.cls.ans.nil} &\xrightarrow{\text{req}} (\overline{\text{ans}}.\text{yes}) \triangleleft \text{cls.ans.nil} && \text{(Good)} \\ &\xrightarrow{\text{cls}} \text{end} \triangleleft \overline{\text{ans}}.nil \xrightarrow{\overline{\text{ans}}} \text{end} \triangleleft \text{nil} \\ m \triangleleft \text{req.cls.ans.nil} &\xrightarrow{\text{req}} (\overline{\text{ans}}.\text{yes}) \triangleleft \text{cls.ans.nil} && \text{(Bad)} \\ &\xrightarrow{\text{cls}} (\overline{\text{ans}}.\text{yes}) \triangleleft \overline{\text{ans}}.nil \xrightarrow{\overline{\text{ans}}} \text{end} \triangleleft \text{nil} \end{aligned}$$

The transition sequence labelled (Good) describes the behaviour according to the instrumentation rules in Figure 4 where, in particular, when the monitor cannot analyse action `cls`, it is terminated with an inconclusive state, which is then retained for the remaining transitions. By contrast, the transition sequence labelled (Bad) describes the behaviour obtained if we had to leave the state of a monitor unaltered when an action cannot be matched. Precisely, in (Bad), the residual monitor ends up matching the subsequent action, $\overline{\text{ans}}$ and reaching an unintended acceptance verdict. ■

It is worth highlighting a few additional aspects of the instrumentation relation in Figure 4. First, note that the instrumentation LTS models the fact that a monitor does *not* have access to the execution graph of a process, and can only analyse the part of the execution graph exhibited by the process at runtime. We illustrate this in the following example.

Example 5 Recall monitor m defined in Example 4 and process p from Example 1. The monitored system may exhibit the following runtime behaviour

$$m \triangleleft p \xrightarrow{\tau} m \triangleleft \text{req}.\overline{\text{ans}}.p + \text{cls}.\text{nil} \xrightarrow{\text{cls}} \text{end} \triangleleft \text{nil}$$

which prohibits the monitor from observing the execution path we saw before at (**) in Example 4, leading to an acceptance verdict. ■

Second, we draw attention to the fact that the instrumentation composition relation is *asymmetric* and lets a monitored system produce a transition as long as a system can produce that transition (even though the monitor *cannot* perform it). This models the requirement that monitors should (as much as possible) have a passive role in an instrumented setting and not interfere with the execution of a system. It also contrasts with the standard intersection and union constructions used in automata theory. We formalise these important properties in Proposition 1 below; for a more comprehensive study of the properties of this instrumentation relation, we refer the reader to [23].

Proposition 1 $m \triangleleft p \xRightarrow{t} m' \triangleleft p'$ iff $p \xRightarrow{t} p'$ and

(i) either $m \xRightarrow{t} m'$

(ii) or $m' = \text{end}$ and $\exists t', \alpha, t'', m'' . t = t'\alpha t''$, $m'' \not\xrightarrow{\tau}$ and $m \xRightarrow{t'} m'' \not\xrightarrow{\alpha}$.

Proof The *only-if* case is proved by numerical induction on the number of transitions n in $m \triangleleft p \xRightarrow{t} m' \triangleleft p'$, where $\xRightarrow{t} = \xrightarrow{\mu_1} \cdot \xrightarrow{\mu_2} \dots \xrightarrow{\mu_n}$ for some $n \geq 0$. The inductive case, i.e., $n = k + 1$, performs a rule case analysis on the first monitored transition. The subcase where rule `ITRM` is used yields subcase (ii).

For the *if* case, we prove the two subcases separately. We first show the first subcase through

$$p \xRightarrow{t} p' \text{ and } m \xRightarrow{t} m' \text{ implies } m \triangleleft p \xRightarrow{t} m' \triangleleft p' \quad (3)$$

by numerical induction on the number of transitions n in $p \xRightarrow{t} p'$, where $\xRightarrow{t} = \xrightarrow{\mu_1} \cdot \xrightarrow{\mu_2} \dots \xrightarrow{\mu_n}$ for some $n \geq 0$. The second subcase is proved by numerical induction on the number of transitions n in $m \xRightarrow{t'} m'' \not\xrightarrow{\alpha}$ and $m'' \not\xrightarrow{\tau}$, where $\xRightarrow{t} = \xrightarrow{\mu_1} \cdot \xrightarrow{\mu_2}$

$\dots \xrightarrow{\mu_2}$ for some $n \geq 0$, and where $t = t'\alpha t''$ for some t', α, t'', m'' , i.e., t can be decomposed into three parts where t' is the prefix of t that can be followed by the monitor, α is the first action that the monitor cannot follow, and t'' is the remaining trace after α . The subproof relies on the property that for every trace t and verdict v , we have $v \xrightarrow{t} v$, and uses property (3) as well. \square

Remark 1 Since we strive towards a general theory of monitorability, the monitor syntax and semantics in Figure 4 allows us to specify non-deterministic monitors such as $\alpha.\text{yes} + \alpha.\text{no}$ or $\alpha.\text{nil} + \alpha.\beta.\text{yes}$. There are settings where determinism is unattainable (e.g., distributed monitoring [25]) or desirable/intended (e.g., in the case of testers [34], or as a means of introducing controlled monitor imperfections to enlarge the set of accepted/analysed traces [36]), and others where non-determinism may be used to express behavioural under-specification, which can be then employed in frameworks for program refinement [1,23]. Thus, expressing non-determinism allows us to study the cases where it is either considered useful, tolerated or considered erroneous. This also allows us to differentiate between different forms of non-deterministic monitor behaviour, as we discuss in more detail in the following sections. \blacksquare

We conclude the section by considering monitors that can detect arbitrarily long execution traces, some of which also behave non-deterministically. We revisit these monitors in subsequent sections.

Example 6 Monitor m_1 (defined below) monitors for runtime executions whose trace is of the form $(\text{req.ans})^*.\text{cls}$ and returns the acceptance verdict **yes**, whereas m_2 dually rejects executions of that form. When composed with process p from Example 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, $m_1 \triangleleft p \xrightarrow{\text{cls}} \text{yes} \triangleleft \text{nil}$.

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

Monitor m_3 may either behave like m_1 or non-deterministically terminate upon a serviced request, i.e., $m_3 \triangleleft p \xrightarrow{\text{req}} \text{req}.m_3 \triangleleft p \xrightarrow{\text{ans}} \text{end} \triangleleft p$. 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 $m_4 \triangleleft p \xrightarrow{\text{cls}} \text{no} \triangleleft \text{nil}$. \blacksquare

5 Correspondence

Although the techniques, algorithms and tools for statically checking system correctness with respect to μHML are well established [3,5], there are situations where checking for system correctness at runtime through monitors is appealing and/or useful. Indeed, the aforementioned static techniques are prone to state-explosion problems in the setting of finite-state systems and are often not applicable to analyse infinite-state systems. In these cases, RV might offer a scalable compromise

for determining system correctness. Alternatively, the system may not be available for analysis prior to deployment, and RV would constitute the only means of analysing system behaviour for correctness.

Our goal is therefore to establish a correspondence between the verdicts reached by monitors over an instrumented system from Section 4 and the properties specified using the logic of Section 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 with respect to some closed μ HML formula, φ . This will, in turn, allow us to determine when a monitor m *represents* (in some precise sense) a property φ .

Example 7 Recall monitor m_1 from Example 6

$$m_1 = \text{rec } x.(\text{req}.\overline{\text{ans}}.x + \text{cls}.\text{yes})$$

Intuitively, one can argue that m_1 monitors for any *satisfactions* of the property

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

This property describes processes that are able to perform a `cls` action after a number of serviced requests, *i.e.*, sequences of `req. $\overline{\text{ans}}$` along one particular execution path. In fact, monitor m_1 produces a `yes` verdict for a computation labelled with a trace in the language $(\text{req}.\overline{\text{ans}})^*.\text{cls}$ from a process p , attesting that $p \in \llbracket \varphi_7 \rrbracket$. Similarly, one can argue that m_2 from Example 6 monitors for *violations* of the property φ_4 from Example 3. The same *cannot* be said for

$$m_4 = \text{rec } x.(\text{req}.\overline{\text{ans}}.x + \text{cls}.\text{yes} + \text{cls}.\text{no})$$

again from Example 6, and φ_7 above. This is because, for some processes, *e.g.*, p from Example 1, m_4 may produce both verdicts `yes` and `no` for the same witness execution traces in $(\text{req}.\overline{\text{ans}})^*.\text{cls}$; this, in turn, leads to contradictions at a logical level since we cannot have both semantic judgements $p \in \llbracket \varphi_7 \rrbracket$ and $p \notin \llbracket \varphi_7 \rrbracket$. A similar argument applies to m_4 and m_2 from Example 6. ■

When instrumented with a process, a monitor may behave non-deterministically in other ways from the behaviour discussed in Example 7 for the case of m_4 .

Example 8 Since the process drives the execution in a monitored system, non-deterministic behaviour from the side of the process may affect the detection capabilities of the instrumented monitor. For instance, we argued in Example 7 that m_1 monitors for property φ_7 . When monitor m_1 is instrumented with process q from Example 1, m_1 may reach an acceptance verdict along various executions such as in the case of $m_1 \triangleleft q \xrightarrow{\text{cls}} \text{yes} \triangleleft \text{nil}$. However, the monitored system may also exhibit the execution

$$m_1 \triangleleft q \xrightarrow{\text{req}.\overline{\text{ans}}} m_1 \triangleleft q_4 \quad \text{where } q_4 = \text{rec } y.\text{req}.\overline{\text{ans}}.y$$

from which m_1 cannot issue an acceptance verdict for q , even though $q \in \llbracket \varphi_7 \rrbracket$. ■

Example 9 Recall $m_3 = \text{rec } x.(\text{req}.\overline{\text{ans}}.x + \text{cls}.\text{yes} + \text{req}.\text{req}.x)$ from Example 6. This monitor may sometimes flag an acceptance but also refrain from doing so when monitoring p from Example 1. This inconsistent behaviour can be observed *even when* p produces the same trace. *E.g.*, for $t = \text{req}.\overline{\text{ans}}.\text{cls}$ we have $m_3 \triangleleft p \xrightarrow{t} \text{yes} \triangleleft \text{nil}$ but also $m_3 \triangleleft p \xrightarrow{t} \text{end} \triangleleft \text{nil}$. ■

There is however a key difference between the non-determinism exhibited by monitor m_4 in Example 7 and that discussed in Example 8 and Example 9. Specifically, in the latter two cases, the non-deterministic behaviour was between detections (*i.e.*, yes and no) and non-detections (*i.e.*, end or any other non-verdict state), which does not lead to any inconsistencies from a logical perspective. By contrast, the non-determinism exhibited by monitor m_4 in Example 7 leads to logical contradictions. This difference provides a blueprint for the subsequent formal framework. Specifically, we investigate conditions that one could require for establishing the correspondence between monitors and formulae. We start with Definition 2, which defines the criteria expected of a monitor m when it *monitors soundly for a property* φ , predicate $\mathbf{smon}(m, \varphi)$ in Definition 2 below, by requiring that acceptances (respectively rejections) imply satisfactions (respectively violations) for every monitored execution of a process p .

Definition 1 (Acceptance and Rejection Predicates)

1. $\mathbf{acc}(p, m) \stackrel{\text{def}}{=} \exists t, p'. m \triangleleft p \xrightarrow{t} \text{yes} \triangleleft p'$
2. $\mathbf{rej}(p, m) \stackrel{\text{def}}{=} \exists t, p'. m \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$ ■

Definition 2 (Sound Monitoring) We say that a monitor m monitors soundly for the property represented by the formula φ , denoted as $\mathbf{smon}(m, \varphi)$, whenever for all processes $p \in \text{PROC}$ the following conditions hold:

1. $\mathbf{acc}(p, m)$ implies $p \in \llbracket \varphi \rrbracket$
2. $\mathbf{rej}(p, m)$ implies $p \notin \llbracket \varphi \rrbracket$ ■

We note that Definition 2 rules out contradicting verdicts by monitors. Whenever $\mathbf{smon}(m, \varphi)$ holds for some monitor m and formula φ , and there exists some process p where $\mathbf{acc}(p, m)$, it must be the case that $p \in \llbracket \varphi \rrbracket$ by Definition 2. Thus, from the logical satisfaction definition we have $\neg(p \notin \llbracket \varphi \rrbracket)$ and by the contrapositive of Definition 2, we must also have $\neg\mathbf{rej}(p, m)$. A symmetric argument also applies for any process p where $\mathbf{rej}(p, m)$, from which $\neg\mathbf{acc}(p, m)$ follows.

Example 10 From Example 7 and Example 9, we can formally state and prove that $\mathbf{smon}(m_1, \varphi_7)$ and $\mathbf{smon}(m_3, \varphi_7)$. Moreover, from Example 3 and Example 6 we can also show that $\mathbf{smon}(m_2, \varphi_4)$.

For instance, recall that $m_2 = \text{rec } x. (\text{req}.\overline{\text{ans}}.x + \text{cls}.\text{no})$, which means that the only verdicts corresponding to logical satisfactions or logical violations that this monitor can produce are rejections. Moreover, from the semantics of Figure 4, each time a rejection is produced by m_2 , this can only be the result of analysing an execution trace in the language $(\text{req}.\overline{\text{ans}})^*.\text{cls}$. But then, from the logic semantics of Figure 3, any such execution trace is enough for the process to violate property $\varphi_4 = \max X. ([\text{req}][\overline{\text{ans}}]X \wedge [\text{cls}]\text{ff})$.

Using Definition 2, we can also show that certain monitors do not soundly monitor for specific formulas, such as $\neg\mathbf{smon}(m_4, \varphi_7)$ and $\neg\mathbf{smon}(m_4, \varphi_4)$. ■

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 Definition 2, *i.e.*, *complete monitoring* for m and φ . Intuitively, this would state 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 Figure 3.

Example 11 Consider the core basic formula $\langle\alpha\rangle\text{tt}$, describing processes that can perform action α . One could check 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, for this same formula, there does *not* exist a sound monitor satisfying $p \notin \llbracket\varphi\rrbracket$ implies $\mathbf{rej}(p, m)$ for any arbitrary process p . We can show this by contradiction. Assume that one such (sound) monitor m exists. Since $\text{nil} \notin \llbracket\langle\alpha\rangle\text{tt}\rrbracket$ then we should have $\mathbf{rej}(\text{nil}, m)$. By Definition 1 and Proposition 1, this means that this particular monitor can produce the behaviour $m \Longrightarrow \text{no}$ which, in turn, implies that $\mathbf{rej}(\alpha.\text{nil}, m)$ also holds although, clearly, $\alpha.\text{nil} \in \llbracket\langle\alpha\rangle\text{tt}\rrbracket$. This makes m unsound, contradicting our initial assumption.

A dual argument can be carried out for another core basic formula in μHML , namely $[\alpha]\text{ff}$, describing the property of not being able to produce action α . 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 $\mathbf{acc}(p, m)$. ■

The core formulas discussed in Example 11 are enough evidence to convince us that requiring complete monitoring would limit correspondence to a trivial subset of the logic. In fact, it would limit it to the semantic subset tt and ff . In view of this, we define weaker forms of completeness, as stated below, where we require completeness with respect to either logical satisfactions or violations (but not both).

Definition 3 (Satisfaction, Violation, and Partially-Complete Monitoring)

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

Using the partially-complete monitoring predicate $\mathbf{cmon}(m, \varphi)$ of Definition 3 and the sound monitoring predicate of Definition 2, we can formalise our touchstone notion for monitor-formula correspondence. Specifically, Definition 4 below states that m *monitors* for formula 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 as stated in Definition 3.

Definition 4 (Monitor-Formula Correspondence)

$$\mathbf{mon}(m, \varphi) \stackrel{\text{def}}{=} \mathbf{smon}(m, \varphi) \text{ and } \mathbf{cmon}(m, \varphi) \quad \blacksquare$$

6 Monitorability

At a general level, *monitorability* is a property of a correctness specification describing the ability to be adequately analysed at runtime; it can also be lifted in pointwise fashion to sets of such specifications. This definition is thus fundamentally dependent on the monitoring setup assumed and the conditions that constitute an adequate runtime analysis. In this section we investigate the monitorability of our logic μHML using Definition 4 as our base notion for adequate monitoring.

Definition 5 (Monitorability) Formula φ is *monitorable* iff there exists a monitor m such that $\mathbf{mon}(m, \varphi)$. A logical 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. The following example gives evidence of this.

Example 12 Through the witness monitor outlined in Example 11, we can show that formulae $\langle\alpha\rangle\text{tt}$ and $\langle\beta\rangle\text{tt}$ are monitorable with a satisfaction complete monitor. However, the composite formula φ_8 defined below is *not*.

$$\varphi_8 = \langle\alpha\rangle\text{tt} \wedge \langle\beta\rangle\text{tt}$$

At an intuitive level, this is the case because once a monitor observes one of the actions required to check either sub-formula $\langle\alpha\rangle\text{tt}$ or $\langle\beta\rangle\text{tt}$, it cannot “go back in time” to observe another action and thus check for the other sub-formula.

Formally, we show this by arguing towards a contradiction. Assume that there exists some monitor m such that $\mathbf{mon}(m, \varphi_8)$. There are two sub-cases to consider:

- If m is satisfaction complete, then $\mathbf{acc}(\alpha.\text{nil} + \beta.\text{nil}, m)$ since $\alpha.\text{nil} + \beta.\text{nil} \in \llbracket \varphi_8 \rrbracket$. By Proposition 1, m reaches verdict **yes** along one of the traces ϵ, α or β . If the trace is the empty trace ϵ , then m must also accept **nil**, which is unsound (since $\text{nil} \notin \llbracket \varphi_8 \rrbracket$). If the trace is α , m must also accept the process $\alpha.\text{nil}$, which is also unsound ($\alpha.\text{nil} \notin \llbracket \varphi_8 \rrbracket$); the case for β is analogous.
- If m is violation complete then $\mathbf{rej}(\beta.\text{nil}, m)$ since $\beta.\text{nil} \notin \llbracket \varphi_8 \rrbracket$. By Proposition 1, we either have $m \xrightarrow{\epsilon} \text{no}$ or $m \xrightarrow{\beta} \text{no}$ and for both cases we can argue that m also rejects process $\alpha.\text{nil} + \beta.\text{nil}$, which is unsound since $\alpha.\text{nil} + \beta.\text{nil} \in \llbracket \varphi_8 \rrbracket$. ■

We now identify the following 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, logical language consists of the safe and co-safe syntactic subsets of μHML , denoted as sHML and cHML respectively.

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

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

Example 13 Recall formulas φ_4 from Example 3 and φ_7 from Example 7. We have $\varphi_4 \in \text{sHML} \subseteq \text{mHML}$ and $\varphi_7 \in \text{cHML} \subseteq \text{mHML}$. However, φ_8 from Example 12 is not in mHML ; this can be easily determined via a simple syntactic analysis. ■

Remark 2 Strictly speaking, the respective definitions for sHML and cHML can be tightened even further. In the case of sHML , we note that **tt** may be expressed as $\max X.X$ since $\llbracket \text{tt} \rrbracket = \llbracket \max X.X \rrbracket$. Dually, in the case of cHML , the formula $\min X.X$ may be written instead of **ff** because $\llbracket \text{ff} \rrbracket = \llbracket \min X.X \rrbracket$. This is in fact the approach adopted in our of earlier works [26]. We decided to include these basic formulas in Definition 6 for a number of reasons. First, **tt** (respectively **ff**) is always considered to be a trivial case of a safety (respectively co-safety) property and omitting it would seem to contradict our target of identifying the maximal

monitorable subsets of the logic μHML (see Section 7); the fact that these basic formulas can be encoded is not immediately discernible from the semantics of Figure 3. Second, having an explicit formula for truth and falsehood is handy for tool implementations and extensions based on the logic. For instance, in [14, 15] a branching construct for reasoning about event data is introduced, and the aforementioned basic formulas facilitate the specification of no-op branches. ■

Clearly $\text{mHML} \subseteq \mu\text{HML}$. To show that mHML is monitorable in the sense of Definition 5, we still need to find a representative monitor for every formula in mHML . Thus, we define a monitor synthesis function $\llbracket - \rrbracket$ generating a monitor for each $\psi \in \text{mHML}$ and then show that $\llbracket \psi \rrbracket$ is the witness monitor required by Definition 5 to demonstrate the monitorability of ψ .

Definition 7 (Monitor Synthesis)

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

A few comments are in order regarding Definition 7. We first note that the synthesis function $\llbracket - \rrbracket$ is *compositional*; see e.g., [34] for reasons why this is desirable. It also assumes a bijective mapping between the denumerable sets $X \in \text{LVAR}$ and $x \in \text{VARS}$; see synthesis for X , $\max X.\psi$ and $\min X.\psi$, where the logical variable X is converted to the process variable x . We also note that, even though Definition 7 covers both sHML and cHML , the syntactic constraints of Definition 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). Finally, we note that the conditional cases used in synthesis of conjunctions, disjunctions, necessity formulas, possibility formulas, and maximal and minimal fixpoints are necessary to be able to generate sound and complete monitors, and handle logically equivalent formulae. We illustrate this in the next two examples.

Example 14 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 of Definition 7 would generate the monitor $\alpha.\text{yes} + \text{no}$ which is not only redundant, but *unsound*. For instance, for process $p = \alpha.\text{nil}$ 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. Consider the cHML formula

$$\langle \alpha \rangle \text{tt} \vee \langle \beta \rangle \text{ff} \vee (\min X. \langle \alpha \rangle \text{ff}) \vee (\langle \alpha \rangle \min X. \text{ff})$$

which also turns out to be logically equivalent to $\langle\alpha\rangle\text{tt}$ because $\llbracket\langle\beta\rangle\text{ff}\rrbracket = \llbracket\text{ff}\rrbracket$, as is $\llbracket\min X.\langle\alpha\rangle\text{ff}\rrbracket = \llbracket\text{ff}\rrbracket = \llbracket\langle\alpha\rangle\min X.\text{ff}\rrbracket$. However, in all of these cases, the monitor synthesis of Definition 7 generates the (sound) monitor $\alpha.\text{yes}$ due to its side conditions. ■

Example 15 Consider the cHML formula $\text{tt}\vee\langle\alpha\rangle\text{tt}$. A naive synthesis without the case checks of Definition 7 would generate the monitor $\text{yes} + \alpha.\text{yes}$ which is clearly *not* violation complete, but its *not* satisfaction complete either. In fact, one can easily see that $\text{nil} \in \llbracket\text{tt}\vee\langle\alpha\rangle\text{tt}\rrbracket$ but $\mathbf{acc}(\text{nil}, \text{yes} + \alpha.\text{yes})$ does not hold, because there is no trace t such that $\text{yes} + \alpha.\text{yes} \triangleleft \text{nil} \xrightarrow{t} \text{yes} \triangleleft \text{nil}$. ■

Remark 3 The monitor synthesis algorithm of Definition 7 may generate non-deterministic monitors for certain formulas in mHML. Consider, for example, the formula

$$\varphi_{10} = (\langle\alpha\rangle\text{tt})\vee(\langle\alpha\rangle\langle\beta\rangle\text{tt}) \in \text{cHML}$$

From Definition 7 we have $\llbracket\varphi_{10}\rrbracket = \alpha.\text{yes} + \alpha.\beta.\text{yes}$, where the generated monitor may sometimes refrain from producing an acceptance verdict for satisfying processes. Specifically, for the process $\alpha.\text{nil} \in \llbracket\varphi_{10}\rrbracket$ we have $\mathbf{acc}(\alpha.\text{nil}, (\alpha.\text{yes} + \alpha.\beta.\text{yes}))$ because of the monitored execution $(\alpha.\text{yes} + \alpha.\beta.\text{yes}) \triangleleft (\alpha.\text{nil}) \xrightarrow{\alpha} \text{yes} \triangleleft \text{nil}$. At the same time, we can also have the alternative monitored execution $(\alpha.\text{yes} + \alpha.\beta.\text{yes}) \triangleleft (\alpha.\text{nil}) \xrightarrow{\alpha} \beta.\text{yes} \triangleleft \text{nil}$. Note, however, that the generated monitor still adequately monitors for the formula, *i.e.*, $\mathbf{mon}(\alpha.\text{yes} + \alpha.\beta.\text{yes}, \varphi_{10})$.

Although we could have defined a more sophisticated monitor synthesis algorithm that handles this non-determinism (at the expense of obscuring further its intended meaning and complicating the ensuing proofs), we opted for a simpler translation that yields correct monitors while still preserving a degree of modularity. An alternative solution could also be to use the synthesis algorithm of Definition 7 and massage mHML formulas before performing the translation (*e.g.*, φ_{10} would be translated to the semantically equivalent formula $\langle\alpha\rangle(\text{tt}\vee\langle\beta\rangle\text{tt})$ which would then be synthesised to the monitor $\alpha.\text{yes}$). ■

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

To prove this theorem, we use $\llbracket-\rrbracket$ to generate witnesses and show that for all $\varphi \in \text{mHML}$, $\mathbf{mon}(\llbracket\varphi\rrbracket, \varphi)$ holds. This proof is given in the dedicated Section 6.1. The reader may safely skip this subsection and proceed to Section 7 upon first reading. Before this, however, we discuss the importance of this result.

Theorem 1 provides us with a simple syntactic check to determine whether a formula is monitorable; as shown earlier in Example 12, determining whether a formula is monitorable is in general non-trivial. The proof of Theorem 1 yields also a pleasing side effect: through Definition 7, we obtain an automatic monitor synthesis algorithm that is correct according to Definition 4.

Example 16 Since φ_4 from Example 3 is in mHML, we know it is monitorable. Moreover, we can generate the *correct* monitor $\llbracket\varphi_4\rrbracket = \text{rec } x.(\text{req}.\overline{\text{ans}}.x + \text{cls}.\text{no}) = m_2$ (from Example 6). Using similar reasoning, φ_7 (Example 7) is also monitorable, and a correct monitor for it is m_1 from Example 6. ■

6.1 Proving Monitorability

Using the monitor synthesis $\llbracket - \rrbracket$ of Definition 7, Theorem 1 is proved in two steps. First we show that the witness monitors generated by $\llbracket - \rrbracket$ are sound, according to Definition 2. In the second step, we show that they are partially complete, according to Definition 3.

We start with soundness. Expanding Definition 2 (Monitor Soundness), we need to show that for all $\psi \in \text{MHML}$, and for all $p \in \text{PROC}$, we have:

$$\mathbf{acc}(p, \llbracket \psi \rrbracket) \text{ implies } p \in \llbracket \psi \rrbracket \quad (4)$$

$$\mathbf{rej}(p, \llbracket \psi \rrbracket) \text{ implies } p \notin \llbracket \psi \rrbracket \quad (5)$$

To simplify the technical exposition, we proceed by proving soundness for the two syntactic subsets, sHML and cHML, separately. We here present the proofs for the subset sHML; those for the cHML subset are similar and can be constructed by the interested reader following a similar procedure.

For sHML, property (4) follows from Lemma 4. This lemma, in turn uses a number of auxiliary lemmata stated below. For instance, Lemma 2 states that if the monitor synthesis of Definition 7 generates a **yes** monitor from a formula $\theta \in \text{sHML}$, then it must be the case that this formula is semantically equivalent to **tt**. Lemma 3 states that the only synthesised monitor from the formulas in sHML that can transition to the verdict **yes** is precisely **yes** itself.

Lemma 1 (Verdict Persistence) $v \xRightarrow{t} m$ implies $m = v$.

Proof By induction on the length of the transition sequence $v \xRightarrow{t} m$. □

Lemma 2 $\forall \theta \in \text{sHML}$. $\llbracket \theta \rrbracket = \mathbf{yes}$ implies $\llbracket \theta \rrbracket = \llbracket \mathbf{tt} \rrbracket$.

Proof By structural induction on θ . Some of the main cases are:

$\theta = \theta_1 \wedge \theta_2$: By Definition 7 $\llbracket \theta_1 \wedge \theta_2 \rrbracket = \mathbf{yes}$ only when $\llbracket \theta_1 \rrbracket = \mathbf{yes}$ and $\llbracket \theta_2 \rrbracket = \mathbf{yes}$. By I.H. we obtain $\llbracket \theta_1 \rrbracket = \llbracket \mathbf{tt} \rrbracket$ and $\llbracket \theta_2 \rrbracket = \llbracket \mathbf{tt} \rrbracket$, from which $\llbracket \theta_1 \wedge \theta_2 \rrbracket = \llbracket \mathbf{tt} \rrbracket$ follows.
 $\theta = \max X.\vartheta$: By Definition 7 $\llbracket \max X.\vartheta \rrbracket = \mathbf{yes}$ only when $\llbracket \vartheta \rrbracket = \mathbf{yes}$. By I.H. we obtain $\llbracket \vartheta \rrbracket = \llbracket \mathbf{tt} \rrbracket$ from which the required result follows. □

Lemma 3 $\forall \theta \in \text{sHML}$. $\llbracket \theta \rrbracket \xRightarrow{t} \mathbf{yes}$ implies $\llbracket \theta \rrbracket = \mathbf{yes}$.

Proof By numerical induction on the length of the transition sequence \xRightarrow{t} .

$n = 0$: Immediate.

$n = k + 1$: We have $\llbracket \theta \rrbracket \xrightarrow{\mu} m \xRightarrow{t'} \mathbf{yes}$. We proceed by case analysis of the structure of θ . We here outline one of the main cases:

$\theta = [\alpha]\vartheta$: From Definition 7 we have two sub-cases to consider:

$\llbracket \vartheta \rrbracket = \mathbf{yes}$: By Definition 7, this implies $\llbracket [\alpha]\vartheta \rrbracket = \mathbf{yes}$ as required.

$\llbracket \vartheta \rrbracket \neq \mathbf{yes}$: By Definition 7, this implies $\llbracket [\alpha]\vartheta \rrbracket = \alpha.\llbracket \vartheta \rrbracket$. We can therefore rewrite $\llbracket \theta \rrbracket \xrightarrow{\mu} m \xRightarrow{t'} \mathbf{yes}$ as $\alpha.\llbracket \vartheta \rrbracket \xrightarrow{\alpha} \llbracket \vartheta \rrbracket$ and $\llbracket \vartheta \rrbracket \xRightarrow{t'} \mathbf{yes}$ where μ could have only been α . However, by $\llbracket \vartheta \rrbracket \xRightarrow{t'} \mathbf{yes}$ and I.H. we know that $\llbracket \vartheta \rrbracket = \mathbf{yes}$ which gives us a contradiction. □

Lemma 4 $\forall \theta \in \text{sHML}, p \in \text{PROC}. \mathbf{acc}(p, \langle \theta \rangle)$ implies $p \in \llbracket \theta \rrbracket$.

Proof By expanding the definition of $\mathbf{acc}(p, \langle \theta \rangle)$, and then by Proposition 1, we know that there exists a trace t such that $\langle \theta \rangle \xrightarrow{t} \text{yes}$. By Lemma 3 we know that $\langle \theta \rangle = \text{yes}$ and subsequently, by Lemma 2, we obtain $\llbracket \theta \rrbracket = \llbracket \text{tt} \rrbracket$, from which $p \in \llbracket \theta \rrbracket$ follows trivially. \square

For sHML, property (5) follows from Lemma 7. This lemma, again uses a number of auxiliary lemmata, some of which have already been presented; the rest are stated below.

Lemma 5 $\forall \theta \in \text{sHML}. p, p' \in \text{PROC}. p \xrightarrow{\tau} p'$ and $p' \notin \llbracket \theta \rrbracket$ implies $p \notin \llbracket \theta \rrbracket$.

Proof In [2](Proposition 3.8 and 3.9), it is shown that:

$$\forall \theta \in \text{sHML}, p \in \text{PROC}. p \in \llbracket \theta \rrbracket \text{ iff } (\forall p' \in \text{PROC}. p \xrightarrow{\epsilon} p' \text{ implies } p' \in \llbracket \theta \rrbracket) \quad (6)$$

The required result follows immediately from the contrapositive of (6) above. \square

Lemma 6 $\forall \theta \in \text{sHML}. \langle \theta \rangle = \text{no}$ implies $\llbracket \theta \rrbracket = \llbracket \text{ff} \rrbracket$.

Proof By structural induction on θ . Most cases are immediate, where the only non-trivial case is:

$\theta = \theta_1 \wedge \theta_2$: By Definition 7 $\langle \theta_1 \wedge \theta_2 \rangle = \text{no}$ only when $\langle \theta_1 \rangle = \text{no}$ or $\langle \theta_2 \rangle = \text{no}$. Without loss of generality, assume the former, i.e., $\langle \theta_1 \rangle = \text{no}$. By I.H. we obtain $\llbracket \theta_1 \rrbracket = \llbracket \text{ff} \rrbracket$ and thus by Figure 3 we deduce $\llbracket \theta_1 \wedge \theta_2 \rrbracket = \llbracket \theta_1 \rrbracket \cap \llbracket \theta_2 \rrbracket = \llbracket \text{ff} \rrbracket \cap \llbracket \theta_2 \rrbracket = \emptyset \cap \llbracket \theta_2 \rrbracket = \emptyset = \llbracket \text{ff} \rrbracket$ as required. \square

Lemma 7 $\forall \theta \in \text{sHML}, p \in \text{PROC}. \mathbf{rej}(p, \langle \theta \rangle)$ implies $p \notin \llbracket \theta \rrbracket$.

Proof By expanding $\mathbf{rej}(p, \langle \theta \rangle)$, we know that there exists a trace t such that $\langle \theta \rangle \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$. The proof thus proceeds by numerical induction on the length of the transition sequence \xrightarrow{t} and then by induction on the structure of θ .

$n = 0$: We thus know that $\langle \theta \rangle = \text{no}$. By Lemma 6 we obtain $\llbracket \theta \rrbracket = \llbracket \text{ff} \rrbracket = \emptyset$, from which the required result follows.

$n = k + 1$: We have $\langle \theta \rangle \triangleleft p \xrightarrow{\mu} m'' \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$. The main sub-cases are:

$\theta = \text{tt}$: Thus $\langle \theta \rangle = \text{yes}$. This however contradicts Lemma 1, since it can never be the case that a yes monitor transitions to a monitor that constitutes a no verdict.

$\theta = [\alpha]\vartheta$: By Definition 7, we have two sub-cases. If $\langle \theta \rangle = \text{yes}$ then we get a contradiction as in the case for $\theta = \text{tt}$. The other sub-case is where $\langle \theta \rangle = \alpha.\langle \vartheta \rangle$. We have three sub-cases to consider:

- $\alpha.\langle \vartheta \rangle \triangleleft p \xrightarrow{\tau} \alpha.\langle \vartheta \rangle \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$ because $p \xrightarrow{\tau} p''$. By I.H. and $\alpha.\langle \vartheta \rangle \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$ we obtain $p'' \notin \llbracket [\alpha]\vartheta \rrbracket$, from which $p \notin \llbracket [\alpha]\vartheta \rrbracket$ follows by $p \xrightarrow{\tau} p''$ and Lemma 5.
- $\alpha.\langle \vartheta \rangle \triangleleft p \xrightarrow{\beta} \text{end} \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$. We get a contradiction by $\text{end} \xrightarrow{t'} \text{no}$ and Lemma 1.
- $\alpha.\langle \vartheta \rangle \triangleleft p \xrightarrow{\alpha} \langle \vartheta \rangle \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$ because $p \xrightarrow{\alpha} p''$. By I.H. we know that $p'' \notin \llbracket \vartheta \rrbracket$ and by $p \xrightarrow{\alpha} p''$ we deduce that $p \notin \llbracket [\alpha]\vartheta \rrbracket$.

$\theta = \theta_1 \wedge \theta_2$: By Definition 7, we have four sub-cases:

- $(\theta_1 \wedge \theta_2) = \text{no}$ because $(\theta_1) = \text{no}$ or $(\theta_2) = \text{no}$. By Lemma 6 we obtain $\llbracket \theta \rrbracket = \llbracket \text{ff} \rrbracket = \emptyset$, from which the required result follows.
- $(\theta_1 \wedge \theta_2) = (\theta_1)$ because $(\theta_2) = \text{yes}$. By I.H. we know $p \notin \llbracket \theta_1 \rrbracket$. Moreover, by $(\theta_2) = \text{yes}$ and Lemma 2 we know $\llbracket \theta_2 \rrbracket = \llbracket \text{tt} \rrbracket$, and thus $\llbracket \theta_1 \wedge \theta_2 \rrbracket = \llbracket \theta_1 \rrbracket$. Thus $p \notin \llbracket \theta_1 \wedge \theta_2 \rrbracket$ follows.
- $(\theta_1 \wedge \theta_2) = (\theta_2)$ because $(\theta_1) = \text{yes}$. Analogous.
- $(\theta_1 \wedge \theta_2) = (\theta_1) + (\theta_2)$. We have four further sub-cases to consider:
 - $(\theta_1) + (\theta_2) \triangleleft p \xrightarrow{\tau} (\theta_1) + (\theta_2) \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$ because $p \xrightarrow{\tau} p''$. By I.H. and $(\theta_1) + (\theta_2) \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$ we obtain $p'' \notin \llbracket \theta_1 \wedge \theta_2 \rrbracket$, from which $p \notin \llbracket \theta_1 \wedge \theta_2 \rrbracket$ follows by $p \xrightarrow{\tau} p''$ and Lemma 5.
 - $(\theta_1) + (\theta_2) \triangleleft p \xrightarrow{\beta} \text{end} \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$. We get a contradiction by $\text{end} \xrightarrow{t'} \text{no}$ and Lemma 1.
 - $(\theta_1) + (\theta_2) \triangleleft p \xrightarrow{\mu} m \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$ because $(\theta_1) \triangleleft p \xrightarrow{\mu} m \triangleleft p''$. This means that $t = \mu t'$. We can thus construct the transition sequence $(\theta_1) \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$ from which we obtain, by I.H., $p \notin \llbracket \theta_1 \rrbracket$ and hence $p \notin \llbracket \theta_1 \wedge \theta_2 \rrbracket$.
 - $(\theta_1) + (\theta_2) \triangleleft p \xrightarrow{\mu} m \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$ because $(\theta_2) \triangleleft p \xrightarrow{\mu} m \triangleleft p''$. Analogous.

$\theta = \max X.\vartheta$: By Definition 7, we have two sub-cases. If $(\theta) = \text{yes}$ then we get a contradiction as in the case for $\theta = \text{tt}$. The other sub-case is where $(\theta) = \text{rec } x.(\vartheta)$. Again we have two sub-cases to consider:

- $\text{rec } x.(\vartheta) \triangleleft p \xrightarrow{\tau} \text{rec } x.(\vartheta) \triangleleft p'' \xrightarrow{t'} \text{no} \triangleleft p'$ because $p \xrightarrow{\tau} p''$. It follows I.H. and Lemma 5, as in the cases above.
- $\text{rec } x.(\vartheta) \triangleleft p \xrightarrow{\tau} (\vartheta)[\text{rec } x.(\vartheta)/x] \triangleleft p \xrightarrow{t'} \text{no} \triangleleft p'$. By I.H. we obtain $p \notin \llbracket \vartheta[\max X.\vartheta/X] \rrbracket$ and, since $\llbracket \max X.\vartheta \rrbracket = \llbracket \vartheta[\max X.\vartheta/X] \rrbracket$, the required result follows. \square

Although we can prove the respectively lemmata for cHML in analogous fashion, we just state them here. From these intermediary results, Proposition 2 follows.

Lemma 8 $\forall \pi \in \text{cHML}, p \in \text{PROC. } \text{acc}(p, (\pi)) \text{ implies } p \in \llbracket \pi \rrbracket$

Lemma 9 $\forall \pi \in \text{cHML}, p \in \text{PROC. } \text{rej}(p, (\pi)) \text{ implies } p \notin \llbracket \pi \rrbracket$

Proposition 2 (Monitorability Soundness) $\forall \psi \in \text{mHML. } \text{smon}((\psi), \psi)$

Proof Immediate from Lemma 4, Lemma 7, Lemma 8 and Lemma 9.

To prove partial completeness of mHML, we again tackle the sub-languages sHML and cHML separately, which allows us to prove finer results. Namely, we can show that all sHML formulae are violation complete with respect to their synthesised monitors of Definition 7, and that all cHML formulae are satisfaction complete. Once again, we provide full details for sHML, and leave the analogous proofs for the cHML case for the interested reader.

Lemma 10 $\forall \theta \in \text{sHML}, p \in \text{PROC. } p \notin \llbracket \theta \rrbracket \text{ implies } \text{rej}(p, (\theta))$

Proof Following [2], the semantic definition of sHML can be given an alternative coinductive characterisation as the largest satisfiability relation satisfying the implications of Table 1 in [2]. In our specific case, in order to prove the required statement, it suffices to show that the relation

$$\mathcal{R} \stackrel{\text{def}}{=} \left\{ (p, \theta) \mid \exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \theta \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p' \right\} \quad (7)$$

is a satisfiability relation. In effect, this proves the contrapositive of the statement we want to show, namely:

$$\neg \text{rej}(p, \llbracket \theta \rrbracket) \text{ implies } p \in \llbracket \theta \rrbracket$$

The proof proceeds by induction on the structure of θ . The main cases are:

$\theta = \theta_1 \wedge \theta_2$: From definition (7), we know that

$$\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \theta_1 \wedge \theta_2 \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'. \quad (8)$$

To show that \mathcal{R} is a satisfiability relation, we have to show that $(p, \theta_1) \in \mathcal{R}$ and $(p, \theta_2) \in \mathcal{R}$. By Definition 7, we have three sub-cases to consider:

- $\llbracket \theta_1 \wedge \theta_2 \rrbracket = \llbracket \theta_1 \rrbracket$ because $\llbracket \theta_2 \rrbracket = \text{yes}$. From (8) we know $\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \theta_1 \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$, hence $(p, \theta_1) \in \mathcal{R}$. Also, since $\llbracket \theta_2 \rrbracket = \text{yes}$, Lemma 1 states that this verdict is persistent. Hence we also know that $\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \theta_2 \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$, which means that $(p, \theta_2) \in \mathcal{R}$ as required.
- $\llbracket \theta_1 \wedge \theta_2 \rrbracket = \llbracket \theta_2 \rrbracket$ because $\llbracket \theta_1 \rrbracket = \text{yes}$. Analogous.
- $\llbracket \theta_1 \wedge \theta_2 \rrbracket = \llbracket \theta_1 \rrbracket + \llbracket \theta_2 \rrbracket$. From (8) we know $\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \theta_1 \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$ and $\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \theta_2 \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$ which imply $(p, \theta_1) \in \mathcal{R}$ and $(p, \theta_2) \in \mathcal{R}$ as required.

$\theta = [\alpha]\vartheta$: From definition (7), we know that

$$\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket [\alpha]\vartheta \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'. \quad (9)$$

According to [2], to show that \mathcal{R} is a satisfiability relation, we have to show the following implication holds: $p \xrightarrow{\alpha} p'$ implies $(p', \vartheta) \in \mathcal{R}$. By Definition 7, we have two sub-cases to consider:

- $\llbracket [\alpha]\vartheta \rrbracket = \text{yes}$ because $\llbracket \vartheta \rrbracket = \text{yes}$. By Lemma 1 we know that, for any p' , $\exists t \in \text{ACT}^*, p'' \in \text{PROC}. \text{yes} \triangleleft p' \xrightarrow{t} \text{no} \triangleleft p''$ which, by $\llbracket \vartheta \rrbracket = \text{yes}$ implies $(p', \vartheta) \in \mathcal{R}$ for any p' . The implication $p \xrightarrow{\alpha} p'$ implies $(p', \vartheta) \in \mathcal{R}$ thus holds trivially.
- $\llbracket [\alpha]\vartheta \rrbracket = \alpha.\llbracket \vartheta \rrbracket$. By (9) we know $\exists t \in \text{ACT}^*, p' \in \text{PROC}. \alpha.\llbracket \vartheta \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$ which either implies:
 - $p \not\xrightarrow{\alpha}$, in which case the implication holds trivially.
 - $\alpha.\llbracket \vartheta \rrbracket \triangleleft p \xrightarrow{\alpha} \llbracket \vartheta \rrbracket \triangleleft p'$ where $p \xrightarrow{\alpha} p'$ and $\exists t' \in \text{ACT}^*, p'' \in \text{PROC}. \llbracket \vartheta \rrbracket \triangleleft p' \xrightarrow{t'} \text{no} \triangleleft p''$, which by I.H., implies that for any such p' , we have $(p', \vartheta) \in \mathcal{R}$ as the implication requires.

$\theta = \max X.\vartheta$: From definition (7), we know that

$$\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \max X.\vartheta \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'. \quad (10)$$

According to [2], to show that \mathcal{R} is a satisfiability relation, we have to show that $(p, \vartheta[\max X.\vartheta/X]) \in \mathcal{R}$. By Definition 7, we have two sub-cases to consider:

- $\llbracket \max X.\vartheta \rrbracket = \text{yes}$ because $\llbracket \vartheta \rrbracket = \text{yes}$. This means that $\llbracket \vartheta[\max X.\vartheta/X] \rrbracket = \text{yes}$ as well, which, by Lemma 1 implies that

$$\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \vartheta[\max X.\vartheta/X] \rrbracket \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'.$$

Hence $(p, \vartheta[\max X.\vartheta/X]) \in \mathcal{R}$ as required.

- $\llbracket \max X.\vartheta \rrbracket = \text{rec } x.\llbracket \vartheta \rrbracket$. From (10) we obtain

$$\exists t \in \text{ACT}^*, p' \in \text{PROC}. \llbracket \vartheta \rrbracket[\llbracket \max X.\vartheta \rrbracket/x] \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'.$$

Hence $(p, \vartheta[\max X.\vartheta/X]) \in \mathcal{R}$ as required. \square

Lemma 11 $\forall \pi \in \text{CHML}, p \in \text{PROC}. p \in \llbracket \pi \rrbracket$ implies $\text{acc}(p, \llbracket \pi \rrbracket)$

Proof For CHML, the satisfaction relation $p \in \llbracket \pi \rrbracket$ can be characterised in standard fashion as the *least* relation \mathcal{R} satisfying the following rules:

- $(p, \text{tt}) \in \mathcal{R}$
- $(p, \pi_1) \in \mathcal{R}$ implies $(p, \pi_1 \vee \pi_2) \in \mathcal{R}$
- $(p, \pi_2) \in \mathcal{R}$ implies $(p, \pi_1 \vee \pi_2) \in \mathcal{R}$
- $p \xrightarrow{\alpha} q$ and $(q, \pi) \in \mathcal{R}$ implies $(p, \langle \alpha \rangle \pi) \in \mathcal{R}$
- $(p, \pi[\min X.\pi/X]) \in \mathcal{R}$ implies $(p, \min X.\pi) \in \mathcal{R}$

The required result thus follows if we show that for any π :

$$(p, \pi) \in \mathcal{R} \quad \text{implies} \quad \text{acc}(p, \llbracket \pi \rrbracket) \quad (11)$$

The proof of property (11) can be using a proof by rule induction on the rules characterising the relation. We here outline the proof for one of the main cases:

- Assume $(p, \langle \alpha \rangle \pi) \in \mathcal{R}$ because

$$p \xrightarrow{\alpha} q \quad (12)$$

$$(q, \pi) \in \mathcal{R} \quad (13)$$

By the I.H., from (13) we know that

$$\text{acc}(q, \llbracket \pi \rrbracket) \quad (14)$$

Recall from Definition 7, $\llbracket \langle \alpha \rangle \pi \rrbracket$ can either be

$$\alpha.\llbracket \pi \rrbracket \quad \text{if} \quad \llbracket \pi \rrbracket \neq \text{no} \quad (15)$$

or no whenever $\llbracket \pi \rrbracket = \text{no}$. Since $\llbracket \pi \rrbracket = \text{no}$ would contradict (14) (see Lemma 1), we rule out the latter case, *i.e.*, $\llbracket \langle \alpha \rangle \pi \rrbracket = \text{no}$, and only consider the case (15).

Now, from (12) and (15) we can deduce that $\llbracket \langle \alpha \rangle \pi \rrbracket \triangleleft p \xrightarrow{\alpha} \llbracket \pi \rrbracket \triangleleft q$ and, as a result of this and of (14), we may conclude $\text{acc}(p, \llbracket \langle \alpha \rangle \pi \rrbracket)$ as required. \square

Proposition 3 (Monitorability Partial Completeness)

$$\forall \psi \in \text{MHML}. \text{cmon}(\llbracket \psi \rrbracket, \psi)$$

Proof Follows from Lemma 10, Lemma 11 and Definition 3. \square

Following these results, Theorem 1 follows.

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

Proof Immediate from Proposition 2 and Proposition 3. \square

7 Expressiveness

The results obtained in Section 6 beg the question of whether μHML 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 μHML according to Definition 4. However, this approach quickly runs into problems since there are monitors, such as $\text{yes} + \text{no}$, that make little sense from the point of view of Definition 4.

To this end, we prove a general (and perhaps surprising) result. Theorem 2 below asserts that multi-verdict monitors are necessarily *unsound*, at least with respect to properties defined over processes (as opposed to logics defined over other domains such as traces, e.g., [18, 12, 22]).

Theorem 2 (Multi-verdict Monitors) *For all monitors $m \in \text{MON}$*

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

Proof By contradiction. Assume that there exists a logical formula $\varphi \in \mu\text{HML}$ such that $\mathbf{mon}(m, \varphi)$. From the traces t and u , we can construct the obvious representative processes for each trace using the construction:

$$\llbracket \alpha.t \rrbracket_{\text{trc2proc}} = \alpha.(\llbracket t \rrbracket_{\text{trc2proc}}) \quad \llbracket \epsilon \rrbracket_{\text{trc2proc}} = \text{nil}$$

We can therefore also construct the process $p = (\llbracket t \rrbracket_{\text{trc2proc}}) + (\llbracket u \rrbracket_{\text{trc2proc}})$, where clearly $p \xrightarrow{t} \text{nil}$ and $p \xrightarrow{u} \text{nil}$. Using the assumptions $m \xrightarrow{t} \text{yes}$, $m \xrightarrow{u} \text{no}$ and Proposition 1 we can thus deduce both $m \triangleleft p \xrightarrow{t} \text{yes} \triangleleft \text{nil}$ 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 (Definition 2), that $p \in \llbracket \varphi \rrbracket$ and $p \notin \llbracket \varphi \rrbracket$. This is clearly a contradiction. \square

This result also implies that, in order to answer the aforementioned question of whether for every sensible monitor m there exists a formula $\varphi \in \mu\text{HML}$ that corresponds to it, that is $\mathbf{mon}(m, \varphi)$, 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 Definition 7 reveals that all the monitors generated using it are, in fact, uni-verdict.

We further restrict the set of what we deem to be sensible monitors in terms of how verdicts are specified. In particular, since monitor verdicts are expected to be definitive, it makes little sense to place them at the top level of an external choice, e.g., $\text{no} + m$ since the meaning of such verdicts is unclear. In fact, a similarly problematic monitor has already been discussed in Example 15. We therefore consider all such monitors as *ill-formed* and restrict our expressivity analysis to *well-formed* monitors, i.e., the complement of ill-formed monitors. It is not hard to see that the monitors generated by Definition 7 are, in fact, all well-formed; this can be show by a simple induction on the structure of the formula.

We partition the set of well-formed, 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 language $\mathcal{L} \subseteq \mu\text{HML}$ is *expressive-complete* with respect to (well-formed) rejection monitors iff for all $m \in \text{RMON}$, there exists some $\varphi \in \mathcal{L}$ such that $\mathbf{mon}(m, \varphi)$. ■

We show that the language sHML (Definition 6) is rejection expressive-complete. We do so with the aid of a mapping function $\langle\langle - \rangle\rangle$ from a rejection monitor to a corresponding formula in sHML defined below. Definition 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 Definition 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 Definition 7.

Definition 9 (Rejection Monitors to sHML Formulae)

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

In the case of rejection monitors, RMON, the monitorability constraint, Definition 4, relating a monitor m with a formula φ can be simplified to the condition below, since the respectively monitors never produce an acceptance.

$$\forall p \in \text{PROC. } \mathbf{rej}(p, m) \quad \text{iff} \quad p \notin \llbracket \varphi \rrbracket \quad (16)$$

The required result, namely Proposition 4, follows from Lemma 12, the proof of which may be skipped upon first reading.

Lemma 12 For all rejection monitors $m \in \text{RMON}$ and arbitrary process $p \in \text{PROC}$ we have $\mathbf{rej}(p, m)$ iff $p \notin \llbracket \langle\langle m \rangle\rangle \rrbracket$

Proof For the *only-if* case we need to show that

$$m \triangleleft p \xrightarrow{t} \mathbf{no} \triangleleft p' \quad \text{implies} \quad p \notin \llbracket \langle\langle m \rangle\rangle \rrbracket$$

The proof proceeds by numerical induction on l , the length of the transition sequence \xrightarrow{t} and then by induction on the structure of m . The main cases are:

$l = 0$: We necessarily have that $m = \mathbf{no}$. By Definition 9 $\langle\langle \mathbf{no} \rangle\rangle = \text{ff}$ from which $p \notin \llbracket \text{ff} \rrbracket = \emptyset$ follows.

$l = k + 1$: We know that $m \triangleleft p \xrightarrow{\mu} m' \triangleleft p'' \xrightarrow{t} \mathbf{no} \triangleleft p'$. The main cases are:

$m = \alpha.n$: From the structure of the monitor, we know

$$\alpha.n \triangleleft p \xrightarrow{\alpha} n \triangleleft p'' \quad \text{where } p \xrightarrow{\alpha} p'' \quad (17)$$

$$n \triangleleft p'' \xrightarrow{t'} \mathbf{no} \triangleleft p' \quad (18)$$

By (18) and I.H. we know $p'' \notin \llbracket \langle\langle n \rangle\rangle \rrbracket$ and by $p \xrightarrow{\alpha} p''$ from (17) we can conclude that $p \notin \llbracket [\alpha]\langle\langle n \rangle\rangle \rrbracket = \llbracket \langle\langle \alpha.n \rangle\rangle \rrbracket$.

$m = n_1 + n_2$: From the structure of the monitor, we know that we have

$$\text{either } n_1 \triangleleft p \xrightarrow{t} \text{no} \triangleleft p' \quad \text{or } n_2 \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$$

If $n_1 \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$, then by I.H. we obtain $p \notin \llbracket \langle n_1 \rangle \rrbracket$, which implies $p \notin \llbracket \langle n_1 \rangle \wedge \langle n_2 \rangle \rrbracket = \llbracket \langle n_1 + n_2 \rangle \rrbracket$. The sub-case for $n_2 \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$ is analogous.

$m = \text{rec } x.n$: From the structure of the monitor, we know

$$\text{rec } x.n \triangleleft p \xrightarrow{\tau} n[\text{rec } x.n/x] \triangleleft p'' \quad \text{where } p \Longrightarrow p'' \quad (19)$$

$$n[\text{rec } x.n/x] \triangleleft p'' \xrightarrow{t} \text{no} \triangleleft p' \quad (20)$$

In particular, for (19), the instrumentation rules of Figure 4 prohibit process p from performing any visible actions until at least $\text{rec } x.n$ unfolds (using a τ -transition). From (20) and $p \Longrightarrow p''$ of (19) we can construct the transition sequence

$$n[\text{rec } x.n/x] \triangleleft p \xrightarrow{t} \text{no} \triangleleft p' \quad (21)$$

where (21) uses (strictly) one less transition than $\text{rec } x.n \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$. Hence, by the I.H. we obtain $p \notin \llbracket \langle n[\text{rec } x.n/x] \rangle \rrbracket = \llbracket \langle n \rangle [\max X. \langle n \rangle / X] \rrbracket$. Now, by Definition 9 and $\langle \text{rec } x.n \rangle = \max X. \langle n \rangle$, the required result follows from the fact that $\llbracket \max X. \langle n \rangle \rrbracket = \llbracket \langle n \rangle [\max X. \langle n \rangle / X] \rrbracket$.

For the *if* case we need to show that

$$p \notin \llbracket \langle m \rangle \rrbracket \quad \text{implies} \quad m \triangleleft p \xrightarrow{t} \text{no} \triangleleft p' \quad \text{for some } t$$

We exploit the fact that our monitors are all assumed to be guarded, and prove the above statement by induction on the structure of m and the number of consecutive top-level $\text{rec } x.m$ constructs a monitor has. We outline the main cases below:

$m = \alpha.n$: By Definition 9 we know $p \notin \llbracket \langle \alpha \rangle \rrbracket$, which implies that

$$p \xrightarrow{\alpha} p' \quad \text{for some } p' \quad (22)$$

$$p' \notin \llbracket \langle n \rangle \rrbracket \quad (23)$$

By (23) and I.H. we obtain $n \triangleleft p' \xrightarrow{t'} \text{no} \triangleleft p''$ for some t', p'' . Thus by $p \xrightarrow{\alpha} p'$ from (22) we derive $\alpha.n \triangleleft p \xrightarrow{\alpha t'} \text{no} \triangleleft p''$ as required.

$m = \text{rec } x.n$: By Definition 9 we know $p \notin \llbracket \max X. \langle n \rangle \rrbracket = \llbracket \langle n \rangle [\max X. \langle n \rangle / X] \rrbracket$. Note that $\langle n[\text{rec } x.n/x] \rangle = \langle n \rangle [\max X. \langle n \rangle / X]$, and since $n[\text{rec } x.n/x]$ must have fewer top-level consecutive $\text{rec } y.-$ constructs than $\text{rec } x.n$ (processes and monitors are assumed to be guarded), we can use $p \notin \llbracket \langle n \rangle [\max X. \langle n \rangle / X] \rrbracket$ and I.H. to obtain $n[\text{rec } x.n/x] \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$ for some t, p' . From this, we can derive $\text{rec } x.n \triangleleft p \xrightarrow{t} \text{no} \triangleleft p'$ as required. \square

Proposition 4 *sHML is Rejection Expressive-Complete.*

Proof Follows from Lemma 12, (16), and the fact that the co-domain of $\langle - \rangle$ is that of sHML formulae. \square

Definition 10 (Acceptance Expressive-Complete) Language $\mathcal{L} \subseteq \mu\text{HML}$ is *expressive-complete* with respect to acceptance monitors iff for all $m \in \text{AMON}$ there exists some $\varphi \in \mathcal{L}$ such that $\mathbf{mon}(m, \varphi)$. ■

Proposition 5 cHML is *Acceptance Expressive-Complete*.

Proof The proof is analogous to that of Proposition 4. □

Equipped with Proposition 4 and Proposition 5, it follows that mHML is expressive complete with respect to uni-verdict monitors.

Definition 11 (Expressive-Complete) $\mathcal{L} \subseteq \mu\text{HML}$ is *expressive-complete* with respect to well-formed uni-verdict monitors iff for all $m \in \text{AMON} \cup \text{RMON}$ there exists some $\varphi \in \mathcal{L}$ such that $\mathbf{mon}(m, \varphi)$. ■

Theorem 3 mHML is *Expressive-Complete*.

Proof Follows immediately from Proposition 4 and Proposition 5. □

We are now in a position to prove the result alluded to at the beginning of this section, namely that mHML is the largest monitorable subset of μHML up to logical equivalence, *i.e.*, Theorem 4 to follow. First, however, we define what we understand by language inclusion up to formula semantic equivalence, Definition 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, which gives an upper bound to the expressiveness of languages satisfying monitorability properties.

Proposition 6 For any $\mathcal{L} \subseteq \mu\text{HML}$:

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

$$\exists m \in \text{AMON}. (\forall p. (\mathbf{acc}(p, m) \text{ iff } p \in \llbracket \varphi \rrbracket)). \quad (24)$$

By Proposition 5, for the monitor m used in (24), we also know

$$\exists \pi \in \text{cHML}. (\forall p. (\mathbf{acc}(p, m) \text{ iff } p \in \llbracket \pi \rrbracket)). \quad (25)$$

Assume an arbitrary $p \in \llbracket \varphi \rrbracket$. By (24) we obtain $\mathbf{acc}(p, m)$, and by (25) we obtain $p \in \llbracket \pi \rrbracket$. Thus $\llbracket \varphi \rrbracket \subseteq \llbracket \pi \rrbracket$. Dually, we can also reason that $\llbracket \pi \rrbracket \subseteq \llbracket \varphi \rrbracket$. □

Theorem 4 (Completeness) If a language $\mathcal{L} \subseteq \mu\text{HML}$ is monitorable then \mathcal{L} cannot (semantically) express more properties than mHML , that is $\mathcal{L} \sqsubseteq \text{mHML}$.

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

$$\forall \varphi \in \mathcal{L}. \exists m \text{ such that } \mathbf{smon}(m, \varphi) \text{ and } \mathbf{cmon}(m, \varphi) \quad (26)$$

By Theorem 2, we know that every m used in (26) 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:

$$(\forall \varphi \in \mathcal{L}_{\text{acc}}. \exists m \in \text{AMON}. \forall p. (\mathbf{acc}(p, m) \text{ iff } p \in \llbracket \varphi \rrbracket)) \quad (27)$$

$$(\forall \varphi \in \mathcal{L}_{\text{rej}}. \exists m \in \text{RMON}. \forall p. (\mathbf{rej}(p, m) \text{ iff } p \notin \llbracket \varphi \rrbracket)) \quad (28)$$

By (27), (28) and Proposition 6 we obtain $\mathcal{L}_{\text{acc}} \sqsubseteq \text{cHML}$ and $\mathcal{L}_{\text{rej}} \sqsubseteq \text{sHML}$ respectively, from which the required result follows. \square

Theorem 2 and Theorem 4 constitute powerful results with respect to the monitorability of the branching-time logic μHML . Completeness, Theorem 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 . 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, Theorem 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 [12], generally carry *higher overheads* than uni-verdict monitors.

Example 17 By virtue of Theorem 4, we can conclude that properties $\varphi_1, \varphi_2, \varphi_3, \varphi_5$ and φ_6 from Example 3 are all non-monitorable properties according to Definition 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 respectively properties; recall that Definition 5 (Monitorability) — through Definition 2 and Definition 3 — universally quantifies over all processes, which generally poses problems for automation.

For instance, in the case of φ_1 , we could use Theorem 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 Example 11). 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 (Definition 3) it should reject the process $\text{req.nil} + \text{cls.nil}$ since this process does not satisfy φ_1 : by Proposition 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 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 $\text{req.ans.nil} + \text{cls.nil}$. Thus, the monitor must be unsound, meaning that it cannot be a rejection monitor. \blacksquare

8 Conclusion

We have investigated monitorability aspects of a branching-time logic called μ HML, a reformulation of the modal μ -calculus that is regarded as one of the most expressive temporal logics [29,30] for graph-based system descriptions (such as LTSs). More specifically, we identified a syntactic subset of this logic, Definition 6, that is the largest monitorable subset in terms of semantic expressivity, Theorem 1 and Theorem 4. We conjecture that our results and methodologies are applicable, at least in part, to other logical specifications that are not necessarily bespoke for RV, thereby extending the utility of the verification technique to a wider class of program logics. Our methods for delineating monitorable logical subsets can also be useful as a foundation for establishing guiding principles when apportioning verification across the pre-deployment and post-deployment phases of software development [4,20].

Future Work: The results obtained in this work are cast with respect to a specific definition of monitorability, Definition 5, which relies on a particular set of requirements for adequate monitoring, Definition 4. However, there are arguably other valid notions of adequate monitoring, depending the application targeted, and it is worth exploring how such alternative definitions for monitorability affect the monitorable subset of μ HML identified in this work. For instance, one could relax the conditions of Definition 4 by only requiring soundness (Definition 2), or require more stringent conditions with respect to verdicts and monitor non-determinism; see [26] for a practical motivation of this.

Monitorability is, in turn, largely dependent on the underlying monitoring setup and instrumentation relation. Potentially, there are other sensible relations apart from the one defined in Figure 4 that are worth considering and it would be interesting to investigate whether this impacts on the monitorable logical subset identified in this work.

In a parallel line of research, we intend to investigate manipulation techniques that decompose formulae into monitorable components as proposed in [20]. For instance, reformulating a generic μ HML formula φ as the disjunction $\phi \vee \varpi$ (recall that $\varpi \in \text{cHML}$) could allow for a *combined* verification approach where ϖ can be adequately verified at runtime using an automatically synthesised dedicated monitor, (ϖ) , whereas the remainder formula ϕ can be verified using traditional techniques such as model checking.

Related Work: We are not the first to address the issue of monitorability in specification logics. In [34], a notion of 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. They however do not investigate any maximally expressive monitorable subset of the logic. The closest result in this regard is that found in [18]: the authors identify an LTL subset for which satisfactions are backed up by derivations in a sound proof system for the logic, and subsequently show how a monitoring algorithm can be automated to infer such proof derivations. They however do not show whether this subset is in some sense maximal in terms of monitorability. Falcone *et al.* [22] revisit the monitorability definition of [34], and extend it to the Safety-Progress property classification, while proposing an alternative definition in terms

of the structure of the recognising Streett Automata of the respectively 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 these different definitions. For instance, in some sense, one can see sHML (respectively cHML) as characterising the set of processes that generate prefix-closed (respectively suffix closed) languages. We however anticipate such a reconciliation effort to be non-trivial, partly due to the difference in the technical machinery used by the respective bodies of work.

The logic μ HML has been previously studied from a linear-time perspective in [2,16], in order to find subsets that characterise may/must testing equivalences. Tests are different from the monitors used in our study (e.g., they actively interact with the system under scrutiny) and the composition relation between tests and systems is distinctly different from the our monitor instrumentation relation. Moreover, their notion of successful computations is disparate from our definition of acceptances and rejections (e.g., in a testing setup, success may be revoked and computations are anonimised, whereas in monitored executions detections are persistent and may be distinguished by their visible trace of actions); see [23] for a more detailed comparison. Interestingly however, the logic subset identified in [2] is related to mHML; in fact the testable subset with respect to may testing coincides with sHML.

Subsets of the logic μ HML have already been used to develop RV tools for checking the behaviour of reactive systems. In [26], the monitorable subset sHML was used to specify safety properties of concurrent Erlang programs. For this logical subset, a monitor synthesis algorithm is defined that generates decentralised monitor arrangements that check for property violations in a concurrent fashion. Based on the work presented in this paper, the tool was then extended to the full monitorable subset in [6], so as to include cHML specifications and acceptance verdicts. In recent work [15,14], the sub-logic sHML has been extended into a form of domain specific language so as to specify adaptation procedures to be taken by the monitor once a violation is detected, thereby mitigating program errors at runtime.

The instrumentation relation used in this work has been investigated further in [23], where it is used to define a number of contextual refinement preorders for monitors that share a similar structure to the ones employed in our study. The work subsequently develops compositional reasoning techniques for the preorders based on our instrumentation relation. In relation to our work, the monitor refinement relations of [23] can be used to compare alternative monitor synthesis procedures such as ours and that of [26] for the monitorable logical subsets identified.

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. W. Ahrendt, J. M. Chimento, G. J. Pace, and G. Schneider. A Specification Language for Static and Runtime Verification of Data and Control properties. In *FM 2015*, volume 9109, pages 108–125. Springer, 2015.

5. J. R. Andersen, N. Andersen, S. Enevoldsen, M. M. Hansen, K. G. Larsen, S. R. Olesen, J. Srba, and J. Wortmann. CAAL: Concurrency Workbench, Aalborg Edition. In *ICTAC*, pages 573–582. Springer, 2015.
6. D. P. Attard and A. Francalanza. A monitoring tool for a branching-time logic. In *RV*, volume 10012 of *LNCS*, pages 473–481. Springer, 2016.
7. C. Baier and J. P. Katoen. *Principles of Model Checking*. MIT Press, New York, May 2008.
8. H. Barringer, Y. Falcone, K. Havelund, G. Reger, and D. E. Rydeheard. Quantified Event Automata: Towards Expressive and Efficient Runtime Monitors. In *FM*, volume 7436 of *LNCS*, pages 68–84. Springer, 2012.
9. H. Barringer, A. Goldberg, K. Havelund, and K. Sen. Rule-Based Runtime Verification. In *VMCAI*, volume 2937 of *LNCS*, pages 44–57. Springer, 2004.
10. 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.
11. A. Bauer, M. Leucker, and C. Schallhart. Comparing LTL semantics for Runtime Verification. *Logic and Computation*, 20(3):651–674, 2010.
12. A. Bauer, M. Leucker, and C. Schallhart. Runtime verification for LTL and TLTL. *Software Engineering and Methodology (TOSEM)*, 20(4):14, 2011.
13. I. Cassar and A. Francalanza. On Synchronous and Asynchronous Monitor Instrumentation for Actor Systems. In *FOCLASA*, volume 175, pages 54–68, 2014.
14. I. Cassar and A. Francalanza. Runtime Adaptation for Actor Systems. In *Runtime Verification (RV)*, volume 9333 of *LNCS*, pages 38–54. Springer, 2015.
15. I. Cassar and A. Francalanza. On Implementing a Monitor-Oriented Programming Framework for Actor Systems. In *iFM*, volume 9681 of *LNCS*, pages 176–192, 2016.
16. A. Cerone and M. Hennessy. Process behaviour: Formulae vs. tests. In *EXPRESS*, volume 41 of *EPTCS*, pages 31–45, 2010.
17. E. Chang, Z. Manna, and A. Pnueli. Characterization of temporal property classes. In *ALP LNCS*, pages 474–486. Springer-Verlag, 1992.
18. C. Cini and A. Francalanza. An LTL Proof System for Runtime Verification. In *TACAS*, volume 9035, pages 581–595. Springer, 2015.
19. E. M. Clarke, Jr., O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, Cambridge, MA, USA, 1999.
20. D. Della Monica and A. Francalanza. Towards a Hybrid Approach to Software Verification. In *NWPT*, number SCS16001 in *RUTR*, pages 51–54. RU Press, 2015.
21. 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.
22. Y. Falcone, J.-C. Fernandez, and L. Mounier. What Can you Verify and Enforce at Runtime? *Software Tools for Technology Transfer (STTT)*, 14(3):349–382, 2012.
23. A. Francalanza. A Theory of Monitors (Extended Abstract). In *FoSSaCS*, volume 9634 of *LNCS*, pages 145–161. Springer, 2016.
24. A. Francalanza, L. Aceto, and A. Ingólfssdóttir. On Verifying Hennessy-Milner Logic with Recursion at Runtime. In *RV*, volume 9333 of *LNCS*, pages 71–86, 2015.
25. A. Francalanza, A. Gauci, and G. J. Pace. Distributed System Contract Monitoring. *Journal of Logic and Algebraic Programming (JLAP)*, 82(5-7):186–215, 2013.
26. A. Francalanza and A. Seychell. Synthesising Correct concurrent Runtime Monitors. *Formal Methods in System Design (FMSD)*, 46(3):226–261, 2015.
27. M. Geilen. On the Construction of Monitors for Temporal Logic Properties. In *RV*, volume 55 of *ENTCS*, pages 181–199, 2001.
28. C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
29. D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science (TCS)*, 27:333–354, 1983.
30. K. G. Larsen. Proof Systems for Satisfiability in Hennessy-Milner Logic with recursion. *Theoretical Computer Science (TCS)*, 72(2):265 – 288, 1990.
31. M. Leucker and C. Schallhart. A Brief Account of Runtime Verification. *Journal of Logic and Algebraic Programming (JLAP)*, 78(5):293 – 303, 2009.
32. Z. Manna and A. Pnueli. Completing the Temporal Picture. *Theoretical Computer Science (TCS)*, 83(1):97–130, 1991.
33. R. Milner. *A Calculus of Communicating Systems*. Springer, 1982.
34. A. Pnueli and A. Zaks. Psl Model Checking and Run-time Verification via Testers. In *FM*, volume 4085, pages 573–586. Springer, 2006.

-
35. K. Sen, G. Rosu, and G. Agha. Generating Optimal Linear Temporal Logic Monitors by Coinduction. In *ASIAN*, volume 2896 of *LNCS*, pages 260–275. Springer, 2003.
 36. J. Shi, S. K. Lahiri, R. Chandra, and G. Challen. Wireless Protocol Validation Under Uncertainty. In *RV*.
 37. detectEr Project. <http://www.cs.um.edu.mt/svrg/Tools/detectEr/>.
 38. A. Vella and A. Francalanza. Preliminary results towards contract monitorability. In *PrePost*, volume 208 of *EPTCS*, pages 54–63, 2016.