Monitoring for Silent Actions*

Luca Aceto[1], Antonis Achilleos[1], Adrian Francalanza[2], and Anna Ingólfsdóttir[1]

1 **School of Computer Science, Reykjavik University, Reykjavik, Iceland**
2 **Dept. of Computer Science, ICT, University of Malta, Msida, Malta**

Abstract

Silent actions are an essential mechanism for system modelling and specification. They are used to abstractly report the occurrence of computation steps without divulging their precise details, thereby enabling the description of important aspects such as the branching structure of a system. Yet, their use rarely features in specification logics used in runtime verification. We study monitorability aspects of a branching-time logic that employs silent actions, identifying which formulas are monitorable for a number of instrumentation setups. We also consider defective instrumentation setups that imprecisely report silent events, and establish monitorability results for tolerating these imperfections.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases Runtime Verification, Monitorability, Hennessy-Milner Logic with Recursion, Silent Actions

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Runtime verification (RV) [17] is a lightweight verification technique that strives to determine whether a system under scrutiny satisfies or violates a property—typically expressed as a formula from some logic—by incrementally analysing its current execution. In general, the runtime analysis is carried out by a *monitor*, a computational entity that observes the exhibited system execution and reaches a verdict once sufficient evidence is observed; the exhibited execution is characterised by a *trace*, a finite sequence of *events* describing the discrete system computational steps. Although the technique may obtain additional (runtime) information that could be useful for verification purposes, it is generally less expressive than exhaustive approaches such as model checking since the verification analysis is limited to the information inferred from the execution trace under consideration. *Monitorability* thus concerns itself with identifying the properties that are analysable by this runtime analysis.

RV setups typically partition computational steps of systems into two groups. On the one hand, *observable* events are those events that are visible (in full) to external entities such as monitors; they are used in the specifications describing system properties and are reported in the system trace. Observable events usually contain runtime data associated with that event (*e.g.* a method-call event would carry information relating the receiver, the method name and the arguments passed as parameters). On the other hand, *unobservable* events broadly encompass the computational steps that are abstracted away either from system

* This research was supported by the project "TheoFoMon: Theoretical Foundations for Monitorability" (grant number: 163406-051) of the Icelandic Research Fund.

modelling or from the respective property specifications; RV setups may occasionally remove these events from a trace so as to allow for a smoother monitoring process [10, 13].

In this work we investigate events that broadly fall somewhere in between these two groups. Concretely, *silent* events (or actions) are computational steps whose specific nature is not disclosed at the level of abstraction of the system model. Nevertheless the model still provides enough evidence of their manifestation during execution, which may play an important role in capturing vital behavioural aspects of the system: they may describe the branching structure of the modelled system behaviour [14, 18] or provide a measure of computational cost and efficiency [4]. In practice, one comes across various instances of such events. For example, the precise details of reported computational steps may be abstracted away for confidentiality/security reasons. Alternatively, the monitoring setup may be unable to report the details of certain computations due to limitations in the instrumentation technology used. In cyber-physical systems, there are also cases where one could detect the occurrence of certain (internal) computation by way of indirect means, such as via the sound of a running motor or the increase in temperature of an enclosed object. For these reasons, behavioural specifications often include descriptions involving silent actions. However, it is unclear how these silent actions are best handled in an RV setup. It is even less clear to what extent silent actions affect the monitorability of the respective specifications.

Our goals are to develop a foundational framework in which these questions may be addressed, and to logically characterize the properties that are monitorable within this framework. Following our work [1, 2, 10–12] and that of others [7, 20], we conduct our investigations in a process-calculus setting, where internal actions have long been studied from both behavioural and specification perspectives. Our study considers a standard labelled-transition-system model that represents silent computational steps as τ-transitions [3, 18], and a variant of the modal μ-calculus [15, 16] with *strong* modal operators that also describe τ-transitions. Our main contributions can be found in the middle sections of the paper:

- Section 3 studies the monitorability of this logic *w.r.t.* a number of monitoring setups that handle τ-actions differently, thus generalising the results obtained in [11, 12].
- Sections 4 and 5 investigate the monitorability of the logic for *imperfect* monitoring setups that obscure aspects of the silent system behaviour expressed by the model, and establish results for tolerating such imperfections.

The appendix collects the proofs of the main technical results.

2 Preliminaries

We assume the following disjoint sets: Act, a (possibly empty) set containing *external* actions, and Sil, a finite set containing *silent* actions. We let α range over Act, δ over Sil, and μ over $\text{Act} \cup \text{Sil}$. A *Labelled Transition System* (LTS) on (Act, Sil) is a triple

$$L = \langle P, (\text{Act}, \text{Sil}), \rightarrow_L \rangle,$$

where P is a nonempty set of system states referred to as *processes* $p, q, \ldots$, and $\rightarrow_L \subseteq P \times (\text{Act} \cup \text{Sil}) \times P$ is a transition relation. We write $p \xrightarrow{\mu}_L q$ instead of $(p, \mu, q) \in \rightarrow_L$ and $p \rightarrow_L q$ if $p \xrightarrow{\delta}_L q$ for some $\delta \in \text{Sil}$. We use $p \stackrel{\mu}{\Rightarrow}_L q$ to mean that, in L, p can derive q using a single μ action and any number of silent actions, that is, $p(\rightarrow_L)^*.\ \xrightarrow{\mu}_L .(\rightarrow_L)^* q$. We distinguish between (general) traces $s = \mu_1.\mu_2.\ldots.\mu_r \in (\text{Act} \cup \text{Sil})^*$ and external traces $t = \alpha_1.\alpha_2.\ldots.\alpha_r \in \text{Act}^*$, and use $p \stackrel{s}{\Rightarrow}_L q$ to mean $p \xrightarrow{\mu_1}_L . \xrightarrow{\mu_2}_L \ldots \xrightarrow{\mu_r}_L q$ and $p \stackrel{t}{\Rightarrow}_L q$ to mean $p \stackrel{\alpha_1}{\Rightarrow}_L . \stackrel{\alpha_2}{\Rightarrow}_L \ldots \stackrel{\alpha_r}{\Rightarrow}_L q$. By $p \xrightarrow{\mu}_L$ we mean that there is some q such that $p \xrightarrow{\mu}_L q$. We occasionally omit the subscript L when it is clear from the context.

▶ **Example 1.** The (standard) regular fragment of *CCS* [18] with grammar:

$$p, q \in \textsc{Proc} ::= \mathtt{nil} \quad | \quad \mu.p \quad | \quad p + q \quad | \quad \mathtt{rec}x.p \quad | \quad x,$$

with x from some countably infinite set of variables, and the transition relation defined as:

$$\textsc{Act}\frac{}{\mu.p \xrightarrow{\mu} p} \qquad \textsc{Rec}\frac{p[\mathtt{rec}x.p/x] \xrightarrow{\mu} q}{\mathtt{rec}x.p \xrightarrow{\mu} q} \qquad \textsc{SelL}\frac{p \xrightarrow{\mu} p'}{p + q \xrightarrow{\mu} p'} \qquad \textsc{SelR}\frac{q \xrightarrow{\mu} q'}{p + q \xrightarrow{\mu} q'}$$

constitutes the LTS $\langle \textsc{Proc}, (\textsc{Act}, \{\tau\}), \rightarrow \rangle$ where τ is the only silent action. ◀

Properties about specific processes may be specified via the logic μHML [16], a reformulation of the modal μ-calculus [15].

▶ **Definition 2.** μHML formulae on $(\textsc{Act}, \textsc{Sil})$ are defined by the grammar:

$$\varphi, \psi \in \mu\text{HML} ::= \mathtt{tt} \quad | \quad \mathtt{ff} \quad | \quad \varphi \wedge \psi \quad | \quad \varphi \vee \psi \quad | \quad \langle\mu\rangle\varphi \quad | \quad [\mu]\varphi \quad | \quad \min X.\varphi \quad | \quad \max X.\varphi \quad | \quad X$$

where X comes from a countably infinite set of logical variables $\textsc{LVar}$. For a given LTS $L = \langle P, (\textsc{Act}, \textsc{Sil}), \rightarrow \rangle$, an environment ρ is a function $\rho : \textsc{LVar} \rightarrow 2^P$. Given an environment ρ, $X \in \textsc{LVar}$, and $S \subseteq P$, $\rho[x \mapsto S]$ denotes the environment where $\rho[X \mapsto S](X) = S$ and $\rho[X \mapsto S](Y) = \rho(Y)$, for all $Y \neq X$. The semantics of a μHML formula φ over an LTS L relative to an environment ρ, denoted as $[\![\varphi, \rho]\!]_L$, is defined as follows:

$$\begin{aligned}
[\![\mathtt{tt}, \rho]\!]_L &= P \qquad [\![\mathtt{ff}, \rho]\!]_L = \emptyset \qquad [\![X, \rho]\!]_L = \rho(X) \\
[\![\varphi_1 \wedge \varphi_2, \rho]\!]_L &= [\![\varphi_1, \rho]\!]_L \cap [\![\varphi_2, \rho]\!]_L \qquad [\![\varphi_1 \vee \varphi_2, \rho]\!]_L = [\![\varphi_1, \rho]\!]_L \cup [\![\varphi_2, \rho]\!]_L \\
[\![[\mu]\varphi, \rho]\!]_L &= \left\{ p \mid \forall q.\ p \xrightarrow{\mu} q \text{ implies } q \in [\![\varphi, \rho]\!]_L \right\} \\
[\![\langle\mu\rangle\varphi, \rho]\!]_L &= \left\{ p \mid \exists q.\ p \xrightarrow{\mu} q \text{ and } q \in [\![\varphi, \rho]\!]_L \right\} \\
[\![\min X.\varphi, \rho]\!]_L &= \bigcap \{ S \mid S \supseteq [\![\varphi, \rho[X \mapsto S]]\!]_L \} \\
[\![\max X.\varphi, \rho]\!]_L &= \bigcup \{ S \mid S \subseteq [\![\varphi, \rho[X \mapsto S]]\!]_L \}
\end{aligned}$$

Two formulae φ and ψ are equivalent, denoted as $\varphi \equiv \psi$, when $[\![\varphi, \rho]\!]_L = [\![\psi, \rho]\!]_L$ for every environment ρ and LTS L. We often consider closed formulae and simply write $[\![\varphi]\!]_L$ for $[\![\varphi, \rho]\!]_L$, as their semantics is independent of ρ. ◀

Let $[\textsc{Sil}]\varphi$ stand for $\bigwedge_{\delta \in \textsc{Sil}} [\delta]\varphi$ and $\langle \textsc{Sil} \rangle \varphi$ for $\bigvee_{\delta \in \textsc{Sil}} \langle \delta \rangle \varphi$. Then, the weak versions of the modalities employed in [1, 11, 12] may be expressed as follows:

$$[[\mu]]\varphi \equiv \max X.([\mu]\varphi \wedge [\textsc{Sil}]X) \qquad \langle\langle\mu\rangle\rangle\varphi \equiv \min X.(\langle\mu\rangle\varphi \vee \langle \textsc{Sil} \rangle X).$$

Readers should consult [3, 16], or more recently [1, 12], for more details on μHML.

3 Monitorability

The logic μHML of Section 2 is very expressive. It is also agnostic of the technique to be employed for verification. This level of generality provides an ideal basis for investigating the interplay between silent actions and the RV technique, and permits us to extend our findings to other specification logics (*e.g.* CTL and CTL* [8] can be encoded in μHML [15]). The property of monitorability, however, fundamentally relies on the monitoring setup considered.

Monitor Semantics

$$\textsc{mRec}\ \frac{m[\mathtt{rec}x.m/x] \xrightarrow{\mu} m'}{\mathtt{rec}x.m \xrightarrow{\mu} m'} \qquad \textsc{mSel}\ \frac{m \xrightarrow{\mu} m'}{m+n \xrightarrow{\mu} m'} \qquad \textsc{mAct}\ \frac{}{\mu.m \xrightarrow{\mu} m} \qquad \textsc{mVrd}\ \frac{}{v \xrightarrow{\mu} v}$$

Instrumentation Semantics

$$\textsc{iMon}\ \frac{p \xrightarrow{\mu}_L q \quad m \xrightarrow{\mu}_M n}{m \triangleleft p \xrightarrow{\mu}_{I(M,L)} n \triangleleft q} \qquad \textsc{iTer}\ \frac{p \xrightarrow{\mu}_L q \quad m \overset{\mu}{\not\longrightarrow}_M}{m \triangleleft p \xrightarrow{\mu}_{I(M,L)} \mathtt{end} \triangleleft q} \qquad \textsc{iAbs}\ \frac{p \xrightarrow{\delta}_L q}{m \triangleleft p \xrightarrow{\delta}_{I(M,L)} m \triangleleft q}$$

where $\mu \in \textsc{Act} \cup \textsc{Sil}$, $v \in \{\mathtt{end}, \mathtt{no}\}$, and $\delta \in \textsc{Sil}$.

Table 1 Behaviour and Instrumentation Rules for Monitored Systems

Monitoring Systems: A *monitoring setup* on $(\textsc{Act}, \textsc{Sil})$ is a triple $S = \langle M, I, L\rangle$, where L is a system LTS on $(\textsc{Act}, \textsc{Sil})$, M is a monitor LTS on $(\textsc{Act}, \textsc{Sil})$, and I is the instrumentation describing how to compose L and M into an LTS, denoted by $I(M, L)$, on $(\textsc{Act}, \textsc{Sil})$. We call the pair (M, I) a *monitoring system* on $(\textsc{Act}, \textsc{Sil})$. For $M = (\textsc{Mon}, (\textsc{Act}, \textsc{Sil}), \rightarrow_M)$, $\textsc{Mon}$ is a set of monitor states (ranged over by m) and $\rightarrow_M$ is the *monitor semantics* described in terms of the behavioural state transitions a monitor takes when it analyses trace events $\mu \in \textsc{Act} \cup \textsc{Sil}$. The states of the composite LTS $I(M, L)$ are written as $m \triangleleft p$, where m is a monitor state and p is a system state; the monitored-system transition relation is here denoted by $\rightarrow_{I(M,L)}$. We focus on *rejection* monitors, *i.e.,* monitors with a designated rejection state $\mathtt{no}$, and hence safety fragments of the logic μHML. However, our arguments apply dually to acceptance monitors and co-safety properties; see [11, 12] for details.

▶ **Definition 3.** Fix a monitoring setup $S = \langle M, I, L\rangle$ on $(\textsc{Act}, \textsc{Sil})$ and let m be a monitor of M and φ a formula of μHML on $(\textsc{Act}, \textsc{Sil})$. We say that m (M, I)*-rejects* (or simply *rejects*, if M, I are evident) a process p in L, written as $\mathbf{rej}_S(m, p)$, when there are a process q in L and a trace $s \in (\textsc{Act} \cup \textsc{Sil})^*$ such that $m \triangleleft p \overset{s}{\Rightarrow}_{I(M,L)} \mathtt{no} \triangleleft q$. We say that m (M, I)*-monitors for* φ on L whenever

$$\text{for each process } p \text{ of } L,\ \mathbf{rej}_S(m, p) \text{ if and only if } p \notin [\![\varphi]\!]_L.$$

Finally, m (M, I)*-monitors for a formula* φ when m (M, I)-monitors for φ on L for every LTS L on $(\textsc{Act}, \textsc{Sil})$. The monitoring system (M, I) is often omitted when evident. ◀

Monitoring for Silent Actions: The first monitoring system we consider does *not* distinguish between silent actions and external actions.

▶ **Definition 4.** A *full monitor* on $(\textsc{Act}, \textsc{Sil})$ is defined by the grammar:

$$m, n \in \textsc{Mon}_\delta ::= \mathtt{end} \quad | \quad \mathtt{no} \quad | \quad \mu.m \quad | \quad m + n \quad | \quad \mathtt{rec}\ x.m \quad | \quad x,$$

where x comes from a countably infinite set of monitor variables. Constant $\mathtt{no}$ denotes the *rejection verdict* state whereas $\mathtt{end}$ denotes the *inconclusive verdict* state. The rules in Table 1 describe the behaviour for full monitors (we elide the obvious symmetric rule for $m + n$). ◀

Note that rule $\textsc{mVrd}$ in Table 1 describes how verdicts are irrevocable; monitors can therefore only describe suffix-closed behaviour.

▶ **Definition 5.** For any system LTS L and monitor LTS M agreeing on $(\textsc{Act}, \textsc{Sil})$, a *full instrumentation* LTS, denoted by $\rightarrow_{I(M,L)}$, is defined by rules $\textsc{iMon}$ and $\textsc{iTer}$ in Table 1. ◀

In rule iMon, when the system produces a trace event μ that the monitor is able to analyse by transitioning from m to n, the constituent components of a monitored system $m \triangleleft p$ move in lockstep. Conversely, when the system produces an event μ that the monitor is *unable* to analyse, the monitored system still executes, according to iTer, but the monitor transitions to the inconclusive state, where it remains for the rest of the computation.

We refer to the monitor LTS in Definition 4 as M^δ, the full instrumentation of Definition 5 as I^δ and the pair (M^δ, I^δ) as the *full monitoring system*. For each system LTS L that agrees with the full monitoring system on (Act, Sil), we can show a correspondence between the respective monitoring setup $\langle M^\delta, I^\delta, L\rangle$ and the following syntactic subset of μHML.

▶ **Definition 6.** The *strong safety* μHML is defined by the grammar:

$$\theta, \chi \in \text{ssHML} ::= \texttt{tt} \quad | \quad \texttt{ff} \quad | \quad [\mu]\theta \quad | \quad \theta \wedge \chi \quad | \quad \max X.\theta \quad | \quad X \qquad \blacktriangleleft$$

As opposed to sHML from [11, 12], ssHML is defined using strong transitions $p \xrightarrow{\mu} q$ (not weak ones, $p \stackrel{\mu}{\Rightarrow} q$) and the modalities $[\mu]\theta$ employ *any* action μ, not just external ones.

▶ **Definition 7.** Fix a monitoring system (M, I), a fragment Λ of μHML, and an LTS L on (Act, Sil). We say that (M, I) monitors for Λ on L whenever:

- For all $\varphi \in \Lambda$, there exists some $m \in M$ that monitors for it on L.
- For all $m \in M$, there exists some $\varphi \in \Lambda$ that is monitored by it on L.

We say that (M, I) monitors for Λ when it monitors for Λ on every LTS L. $\blacktriangleleft$

▶ **Theorem 8.** *The* full monitoring system (M^δ, I^δ) *monitors for* ssHML*.*

Proof. See Appendix A.2. $\blacktriangleleft$

Monitoring for External Actions: The results obtained in [11, 12] can be expressed and recovered within our more general framework.

▶ **Definition 9.** *Safety* μHML, presented in [11, 12], is defined by the grammar:

$$\pi, \kappa \in \text{sHML} ::= \texttt{tt} \quad | \quad \texttt{ff} \quad | \quad [[\alpha]]\pi \quad | \quad \pi \wedge \kappa \quad | \quad \max X.\pi \quad | \quad X.$$

Note that $[[\alpha]]\pi$ uses *external actions*. Its semantics is given as in Definition 2. We can also give a direct inductive definition, *i.e.*, $[\![[[\alpha]]\varphi, \rho]\!] = \{p \mid \forall q.\ p \stackrel{\alpha}{\Rightarrow} q \text{ implies } q \in [\![\varphi, \rho]\!]\}$. $\blacktriangleleft$

▶ **Definition 10.** An *external monitor* on (Act, Sil) is defined by the grammar:

$$m, n \in \text{Mon}_\alpha ::= \texttt{end} \quad | \quad \texttt{no} \quad | \quad \alpha.m \quad | \quad m + n \quad | \quad \texttt{rec}\ x.m \quad | \quad x.$$

Table 1 defines its LTS transition semantics, yielding $M^\alpha = \langle \text{Mon}_\alpha, (\text{Act}, \text{Sil}), \rightarrow\rangle$. *External instrumentation*, denote by I^α, is defined by the *three* rules iMon, iTer, and iAbs in Table 1; in the case of iMon and iTer action μ is substituted by the external action α. We refer to the pair (M^α, I^α) as the *external monitoring system*, amounting to the setup in [11, 12]. $\blacktriangleleft$

▶ **Theorem 11.** *The* external monitoring system (M^α, I^α) *monitors for sublogic* sHML*.* $\blacktriangleleft$

▶ **Example 12.** Consider a simple server interface that receives requests from a client, represented by action `req`, and then sends a reply, represented by action `ans`. Between `req` and `ans`, a server implementation may upload a copy of the request transcript; this computation is represented as a sequence of silent τ-transitions that do not divulge information relating to the upload. In LTS L_1 of Figure 1, process 1 represents a server implementation

Figure 1 LTS L_1 depicts the two variations of the server from Example 12.

that never uploads anything, whereas process 4 represents an alternative implementation that creates a transcript (the τ-transition from 5 to 6) and repeatedly attempts to upload the copy until it succeeds (the τ-loop on 6 followed by the transition to 7). An external monitor does not see processes 1 and 4 differently, as it does not observe the silent transitions. On the other hand, a full monitor can observe all the silent transitions that occur during an execution. We note that both process 1 and process 4 in L_1 violate the specification $[[\texttt{req}]][[\texttt{ans}]]\texttt{ff}$. Process 1 violates $[\texttt{req}][\texttt{ans}]\texttt{ff}$, while 4 does not. Conversely, process 1 does not violate the ssHML-specification $[\texttt{req}][[\tau]][\texttt{ans}]\texttt{ff}$, but 4 does: this can be observed by the full monitor $\texttt{req}.\texttt{rec}\ x.(\tau.(\texttt{ans}.\texttt{no} + x))$. ◀

We conclude the section by commenting on other potential monitoring systems and their expressive power. In particular, the monitoring system (M^δ, I^α) yields monitoring setups whereby monitor δ-transitions are suppressed by the instrumentation, effectively making full monitors behave like external monitors from Definition 10. In the case of the monitoring system (M^α, I^δ), the instrumentation forces the monitor to transition to the inconclusive state more often since it does not abstract away from δ-transitions.

4 Obscuring the Silent Transitions

The full monitoring system (M^δ, I^δ) presented in Section 3 is straightforward and powerful. One might however argue that, in practice, it is *too* powerful: it is plausible that the visibility of certain silent transitions be somehow more *obscure* than that of external transitions. The external monitoring system (M^α, I^α) sits at the other end of the spectrum because it completely ignores all silent transitions. We consider monitoring systems that fall between these extremes: they can clearly observe certain silent transitions, but may receive imperfect information on others *i.e.,* observing that some number of transitions occurred, but *not* how many. In this case, we say that the transitions were *obscure*.

4.1 A Preorder of Obscure LTSs and Reliable Monitoring

We consider two silent actions: τ is a silent action that can be clearly observed and υ is the *obscure* silent action, representing an undetermined positive number of τ-transitions. In the following, we consider only monitoring setups on $(\text{Act}, \{\tau, \upsilon\})$ and, whenever we say that L is an LTS, we mean that it is a system LTS on $(\text{Act}, \{\tau, \upsilon\})$, unless otherwise stated; if L reports perfect information, it is assumed to be an LTS on $(\text{Act}, \{\tau\})$.

We consider a preorder $\leq_o$ on LTSs, where $L \leq_o L'$ intuitively means that L and L' have the same processes, but the silent transitions in L' are somehow more obscure than in L. Although we do not identify a specific such preorder, in Subsection 4.2, we introduce properties that we require of it. We say that L' is *an obscuring* of L when $L \leq_o L'$. We also introduce the obscuring preorder $\leq$ on $\text{Act} \cup \{\tau, \upsilon\}$: $\mu_1 \leq \mu_2$ iff $\mu_1 = \mu_2$ or $\mu_1 = \tau$ and

$\mu_2 = \upsilon$. The intuition is that, whenever $\mu_1 \leq \mu_2$ and the system performs a μ_1-transition, the monitor may observe a (more obscure) μ_2-transition.

▶ **Example 13.** Consider a simple LTS L which contains exactly *one* maximal path:

$$p \xrightarrow{\mu_1}_L p_1 \xrightarrow{\mu_2}_L \cdots \xrightarrow{\mu_i}_L p_i \xrightarrow{\tau}_L q_1 \xrightarrow{\tau}_L \cdots \xrightarrow{\tau}_L q_r \xrightarrow{\mu_{i+1}}_L p_{i+1} \xrightarrow{\mu_{i+2}}_L \cdots \xrightarrow{\mu_k}_L p_k$$

of $k+r$ transitions, where $r > 0$; note that states $q_2 \ldots q_{r-1}$ have no outgoing external actions. An obscuring of L may result from replacing $p_i \xrightarrow{\tau}_L q_1$ by a direct transition $p_i \xrightarrow{\upsilon}_{L'} q_r$, thus obscuring the path $p_i \xrightarrow{\tau}_L q_1 \xrightarrow{\tau}_L \cdots \xrightarrow{\tau}_L q_r$ and leaving the remaining path unchanged. Thus in L' we have:

$$p \xrightarrow{\mu_1}_{L'} p_1 \xrightarrow{\mu_2}_{L'} \cdots \xrightarrow{\mu_i}_{L'} p_i \xrightarrow{\upsilon}_{L'} q_r \xrightarrow{\mu_{i+1}}_{L'} p_{i+1} \xrightarrow{\mu_{i+2}}_{L'} \cdots \xrightarrow{\mu_k}_{L'} p_k$$

This would mean that as the system progresses from p to p_k, μ_1 through μ_k are clearly observed, but when the system performs $p_i \xrightarrow{\tau}_L q_1 \xrightarrow{\tau}_L \cdots \xrightarrow{\tau}_L q_r$, we only observe that at least one silent transition occurred, without discerning the exact number. ◀

▶ **Definition 14.** Let m be a monitor of a monitoring system (M, I); φ a formula of μHML on $(\text{Act}, \{\tau\})$; L an LTS on $(\text{Act}, \{\tau\})$; and L' an obscuring of L. We say that m (M, I)-*monitors* for φ on L from L' iff

$$\text{for every process } p \text{ of } L',\ p \notin [\![\varphi]\!]_L \text{ if and only if } \mathbf{rej}_{\langle M,I,L'\rangle}(m, p).$$

We say that m *reliably* (M, I)*-monitors* for φ on L if m (M, I)-monitors for φ on L from any obscuring of L. Monitor m reliably (M, I)-monitors for φ if m reliably (M, I)-monitors for φ on any LTS L. We often omit the the monitoring system (M, I) whenever it is evident. ◀

▶ **Definition 15.** Fix a monitoring system (M, I) and a fragment Λ of μHML on $(\text{Act}, \{\tau, \upsilon\})$. (M, I) *reliably monitors* for Λ on LTS L iff

- For every $\varphi \in \Lambda$, there is a monitor m of M such that m reliably monitors for φ on L.
- For every a monitor m of M, there is a $\varphi \in \Lambda$ such that m reliably monitors for φ on L.

(M, I) *reliably monitors* for Λ when (M, I) *reliably monitors* for Λ on every LTS. ◀

4.2 Requirements on Obscuring Preorders

We identify certain properties of the obscuring ordering $\leq_o$ that we consider natural. These properties suffice to prove the results of Section 5. Consequently, the conclusions we draw about reliably monitorable formulas of μHML are proven for every $\leq_o$ that has these properties. Our intuition is that if $L \leq_o L'$, then L' is the same LTS as L, but seen with less precision with respect to the silent transitions. So, every transition we observe in L' is either a transition from L, or an obscure view of a sequence of transitions from L.

Natural Properties of Obscurings. We fix two LTSs $L \leq_o L'$. Since L' should at most provide imperfect information on the *silent* transitions of the system, external transitions should be unaffected:

A. $\xrightarrow{\alpha}_{L'} = \xrightarrow{\alpha}_L$ for every $\alpha \in \text{Act}$.
As L' obscures the information on the silent transitions of L, τ-transitions will become fewer: L' should have at most the τ-transitions of L (Property B). Furthermore, every υ-transition in L' represents a non-empty sequence of silent transitions from L (Property C).

B. $\xrightarrow{\tau}_{L'} \subseteq \xrightarrow{\tau}_L$ and

C. $\xrightarrow{\upsilon}_{L'} \subseteq (\xrightarrow{\tau}_{L} \cup \xrightarrow{\upsilon}_{L})^+$.

The following properties ensure that a certain level of information is retained in L'. In particular, if a state has a silent transition in L, it should still have a silent transition in L' (Property D). Moreover, if a state p has a sequence of silent transitions in L that lead to a state q that can perform an external action, then this observation should be preserved in L'. Following Property A, it suffices to require that q is reachable from p in L' via a sequence of silent actions (Property E).

D. For all p if $p \xrightarrow{\tau}_{L}$ or $p \xrightarrow{\upsilon}_{L}$, then $p \xrightarrow{\tau}_{L'}$ or $p \xrightarrow{\upsilon}_{L'}$.

E. For all p, p', if $p(\xrightarrow{\tau}_{L} \cup \xrightarrow{\upsilon}_{L})^+ p' \xrightarrow{\alpha}$ for some $\alpha \in \text{Act}$, then $p(\xrightarrow{\tau}_{L'} \cup \xrightarrow{\upsilon}_{L'})^+ p'$.

The Strength of Obscuring. Properties F, G, and H capture the kind of obscuring ordering considered in this paper. We assume that there is a certain level of obscuring, beyond which adequate monitoring is deemed infeasible. In Property F below, obscuring can reach a point, represented as an LTS L^o, where all the silent-action information is hidden. That is, if $p \xrightarrow{\upsilon}_{L^o} \xrightarrow{\upsilon}_{L^o} p'$, then process p can also perform the more obscure transition $p \xrightarrow{\upsilon}_{L^o} p'$ and furthermore, at no point does L^o reveal any clear τ-transition. We call such an obscuring L^o, as described by Property F, a *total* obscuring.

F. Each L has an obscuring L^o, such that $\xrightarrow{\tau}_{L^o} = \emptyset$ and $\xrightarrow{\upsilon}_{L^o}$ is transitive.

For an LTS L, let L^τ be the LTS on $(\text{Act}, \{\tau\})$ with the same set of processes, so that for $\alpha \in \text{Act}$, $\xrightarrow{\alpha}_{L^\tau} = \xrightarrow{\alpha}_{L}$ and $\xrightarrow{\tau}_{L^\tau} = \xrightarrow{\tau}_{L} \cup \xrightarrow{\upsilon}_{L}$. Property G assures us that we can always obscure any selection of τ-transitions by turning them into υ-transitions, thus "forgetting" how many transitions were taking place at certain points. Property G can also be interpreted to mean that υ-transitions may indeed just represent single τ-transitions.

G. $L^\tau \leq_o L$ for each LTS L.

For the last requirement, we need the following definitions. For a process p in L and a trace $s \in (\text{Act} \cup \{\tau, \upsilon\})^*$, we say that p *represents* s *in* L when s is the only maximal trace that p can produce — that is, when $\forall s'.\big(\exists q.\ p \overset{s'}{\Rightarrow}_{L} q$ iff s' is a prefix of $s\big)$. For a trace $s \in (\text{Act} \cup \{\upsilon\})^*$, we define the *total obscuring* of s, denoted as $o(s)$, as follows: $o(\epsilon) = \epsilon$; $o(\upsilon^k) = \upsilon$ and $o(\upsilon^k \alpha s) = \upsilon\alpha\, o(s)$ for $k > 0$; and $o(\alpha s) = \alpha\, o(s)$. Property H ensures that any sequence of silent transitions can be obscured into an υ-transition at least for some LTSs:

H. for every trace $s \in (\text{Act} \cup \{\upsilon\})^*$, there are LTSs $L \leq_o L'$ and a process p in L, such that p represents s in L and $o(s)$ in L'.

Property H may seem arbitrary, but it is not hard to justify that it is an immediate consequence of our intuition, as depicted in Example 13. Consider a maximal-path LTS L as in Example 13, but with τ-transitions replaced by υ-transitions, such that $s_1 = \mu_1 \cdots \mu_i \in \text{Act}^*$ and $s_2 = \mu_{i+1} \cdots \mu_k \in \text{Act}^*$. Then, p represents $s = s_1 \upsilon^k s_2$ in L and $o(s) = s_1 \upsilon s_2$ in L'.

4.3 An Ordering of Obscurings

We provide a natural instance of an ordering that has all the properties of Subsection 4.2.

▶ **Definition 16.** Relation $\leq_c$ is the transitive closure of $\leq_1$, where for LTSs L_1 and L_2 on $(\text{Act}, \{\tau, \upsilon\})$, $L_1 \leq_1 L_2$ when for every $\alpha \in \text{Act}$, $\xrightarrow{\alpha}_{L_1} = \xrightarrow{\alpha}_{L_2}$ and one of the following holds:

1. $\xrightarrow{\tau}_{L_1} = \xrightarrow{\tau}_{L_2}$ and $\xrightarrow{\upsilon}_{L_1} \subseteq \xrightarrow{\upsilon}_{L_2} \subseteq \xrightarrow{\upsilon}_{L_1} \cup \xrightarrow{\tau}_{L_1}$;
2. $\xrightarrow{\upsilon}_{L_1} = \xrightarrow{\upsilon}_{L_2}$ and $\xrightarrow{\tau}_{L_2} \subseteq \xrightarrow{\tau}_{L_1} \subseteq \xrightarrow{\upsilon}_{L_2} \cup \xrightarrow{\tau}_{L_2}$;

3. $\xrightarrow{\tau}_{L_1} = \xrightarrow{\tau}_{L_2}$ and $\xrightarrow{\upsilon}_{L_1} \subseteq \xrightarrow{\upsilon}_{L_2} \subseteq (\xrightarrow{\upsilon}_{L_1})^+$; or
4. $\xrightarrow{\tau}_{L_1} = \xrightarrow{\tau}_{L_2}$, $\xrightarrow{\upsilon}_{L_2} \subseteq \xrightarrow{\upsilon}_{L_1}$, and for all $p \xrightarrow{\upsilon}_{L_1} p'$, if $p \not\xrightarrow{\upsilon}_{L_2} p'$, then
   $p' \not\xrightarrow{\alpha}$ for all $\alpha \in \text{Act}$,
   $p' \xrightarrow{\upsilon}_{L_1} p''$ or $p' \xrightarrow{\tau}_{L_1} p''$ for some $p'' \neq p'$, and
   $p \xrightarrow{\upsilon}_{L_2} p''$ for every p'' such that $p' \xrightarrow{\upsilon}_{L_1} p''$ or $p' \xrightarrow{\tau}_{L_1} p''$. ◀

The cases presented in Definition 16 give a set of *moves* we can apply to construct a more obscure LTS from a given one. Informally:

1. According to move 1, for any transition $p \xrightarrow{\tau} q$, we can add transition $p \xrightarrow{\upsilon} q$.
2. Following move 1, we can remove transition $p \xrightarrow{\tau} q$.
3. For transitions $p \xrightarrow{\upsilon} p' \xrightarrow{\upsilon} p''$, we can insert a new transition $p \xrightarrow{\upsilon} p''$.
4. For transition $p \xrightarrow{\upsilon} p'$, if move 3 has already been applied to $p \xrightarrow{\upsilon} p' \xrightarrow{\upsilon} p''$ for all possible and at least one $p' \xrightarrow{\delta} p''$, where $p' \neq p''$ and $\delta \in \{\tau, \upsilon\}$, and $p' \not\xrightarrow{\alpha}$ for all $\alpha \in \text{Act}$ then we can remove $p \xrightarrow{\upsilon} p'$.

▶ **Example 17.** We revisit Example 12 of a simple server. The LTS L_2 of Figure 1 presents a maximal obscuring of L_1, according to $\leq_c$. Moves 1 and 2 can replace all τ-transitions by υ-transitions; move 3 can be used to introduce a transition from process 2 directly to process 5; and move 4 can eliminate incoming transitions to process 4, including the self-loop. Thus, the LTS retains the information that the server uploads the transcript to a remote location, but not any information of intermediate steps. We observe that formula $\psi = [\texttt{req}][[\tau]][\texttt{ans}]\texttt{ff}$ is reliably monitorable on L_1 from L_2 by the full monitor $\texttt{req.rec}\ x.(\tau.(\texttt{ans.no} + x) + \upsilon.(\texttt{ans.no} + x))$. ◀

▶ **Proposition 18.** *Relation $\leq_c$ has all the properties listed in Subsection 4.2.* ◀

5 Reliable Monitorability

In this section, we identify a maximal reliably monitorable fragment of μHML — up to logical equivalence — and a monitoring system that monitors for it. The results of this section are relative to any fixed preorder $\leq_o$ that satisfies the properties presented in Subsection 4.2.

▶ **Example 19.** Let $\varphi_1 = [\tau][\alpha]\texttt{ff}$ (*i.e.,* after any τ-action, a process cannot perform an α-action), $\varphi_2 = [\tau]\texttt{ff}$ (*i.e.,* a process cannot perform a τ-action), and $\varphi_3 = \max X.([\tau][\alpha]\texttt{ff} \wedge [\tau]X)$ (*i.e.,* a process cannot perform an α-action after any non-empty sequence of τ-actions). Notice that $\varphi_3 \equiv [[\tau]][\alpha]\texttt{ff}$. Let L_1, L_2 be the LTSs described below, where $L_1 \leq_o L_2$.:

$$L_1: \quad p_0 \xrightarrow{\tau} p_1 \xrightarrow{\tau} p_2 \xrightarrow{\alpha} p_3 \qquad\qquad L_2: \quad p_0 \xrightarrow{\upsilon} p_2 \xrightarrow{\alpha} p_3 \quad \text{and} \quad p_1 \xrightarrow{\upsilon} p_2.$$

L_2 is a $\leq_o$-maximal obscuring of L_1: any LTS L' with $L_2 \leq_o L'$ will have to be exactly L_2 according to Properties B through E. LTSs L_1 and L_2 are really instances of LTSs L and L' from Example 13, resp. Consider L_3 described below:

$$L_3: \quad p_0 \xrightarrow{\tau} p_2 \xrightarrow{\alpha} p_3 \quad \text{and} \quad p_1 \xrightarrow{\tau} p_2$$

where L_2 is also an obscuring of L_3, $L_3 \leq_o L_2$. We observe that φ_1 is not reliably monitorable according to Definition 14: $p_0 \in [\![\varphi_1]\!]_{L_1}$ and $p_0 \notin [\![\varphi_1]\!]_{L_3}$, so a monitor that reliably monitors for φ_1 would need to reject and not reject p_0 in L_2. On the other hand, both φ_2 and φ_3 are reliably monitorable *w.r.t.* $\leq_o$. Let

$$m_2 = \upsilon.\texttt{no} + \tau.\texttt{no} \quad \text{and} \quad m_3 = \texttt{rec}\ x.(\upsilon.\alpha.\texttt{no} + \upsilon.x + \tau.\alpha.\texttt{no} + \tau.x)$$

$$\textsc{iMon}\frac{p \xrightarrow{\mu}_L p' \quad m \xrightarrow{\mu'}_M m' \quad \mu \leq \mu'}{m \triangleleft p \xrightarrow{\mu'}_{I(L,M)} m' \triangleleft p'} \qquad \textsc{iTer}\frac{p \xrightarrow{\mu}_L p' \quad \forall \mu' \geq \mu.\ m \overset{\mu'}{\not\longrightarrow}_M}{m \triangleleft p \xrightarrow{\mu}_{I(L,M)} \mathtt{end} \triangleleft p'}$$

$$\textsc{iTran}\frac{m \triangleleft p \xrightarrow{\upsilon}_{I(L,M)} m' \triangleleft p' \quad p' \xrightarrow{\mu}_L p'' \quad \mu \leq \upsilon}{m \triangleleft p \xrightarrow{\upsilon}_{I(L,M)} m' \triangleleft p''}$$

Table 2 Instrumentation rules for myopic monitors.

be monitors from the full monitoring system (M^δ, I^δ). According to Properties B, C, and D, for all LTSs $L \leq_o L'$, where L is an LTS on $(\textsc{Act}, \{\tau\})$, and every process p in L we have that $p \xrightarrow{\tau}_L$ if and only if $p \xrightarrow{\tau}_{L'}$ or $p \xrightarrow{\upsilon}_{L'}$; therefore, m_2 monitors for φ_2 on L from L'. A process p of L violates φ_3 iff $p(\xrightarrow{\tau}_L)^+ q \xrightarrow{\alpha}_L$ for some q; by Properties A, B, C, $p(\xrightarrow{\tau}_L)^+ q \xrightarrow{\alpha}_L$ iff $p(\xrightarrow{\tau}_{L'} \cup \xrightarrow{\upsilon}_{L'})^+ q \xrightarrow{\alpha}_{L'}$, and therefore, m_3 monitors for φ_3 on L from L'. ◀

We first introduce myopic monitors, that are equivalent to full monitors on total obscurings.

▶ **Definition 20.** A *myopic monitor* on $(\textsc{Act}, \textsc{Sil})$ is defined by the grammar:

$$m, n \in \textsc{Mon}_\upsilon ::= \mathtt{end} \mid \mathtt{no} \mid \alpha.m \mid \upsilon.m \mid m + n \mid \mathtt{rec}\ x.m \mid x.$$

A myopic monitor's LTS semantics is defined by the transition rules in Table 1; the resulting monitor LTS is $M^\upsilon = \langle \textsc{Mon}_\upsilon, (\textsc{Act}, \{\upsilon\}), \rightarrow \rangle$. The instrumentation I^υ of myopic monitors is then defined by the rules in Table 2. ◀

Rules iMon and iTer are similar to those for I^δ. The difference is that when the monitor is expecting a more obscure action (*i.e.,* υ), the instrumentation can pass along a possibly less obscure process action. So, the instrumentation may interpret τ-transitions as υ-transitions. Rule iTran is new, but the intuition behind it is similar: the instrumentation may interpret a (possibly mixed) sequence of τ- and υ-transitions as a single υ-transition, if that is what the monitor was expecting. On total obscurings, myopic monitors behave like full monitors.

▶ **Lemma 21.** *If L is a total obscuring on $(\textsc{Act}, \{\tau, \upsilon\})$, then for every $m \in \textsc{Mon}_\upsilon$ and process p of L,* $\mathbf{rej}_{\langle M^\upsilon, I^\upsilon, L\rangle}(m, p)$ *iff* $\mathbf{rej}_{\langle M^\delta, I^\delta, L\rangle}(m, p)$. ◀

Lemma 21 allows us to further restrict the syntax of myopic monitors while preserving monitorability with respect to reliably monitorable formulas. The *υ-alternating myopic monitors* are the myopic monitors restricted to the following syntax:

$$m, n \in \textsc{Mon}_{alt} ::= \mathtt{end} \mid \mathtt{no} \mid \upsilon.\mathtt{no} \mid \alpha.m \mid \upsilon.\alpha.m \mid m{+}n \mid \mathtt{rec}\ x.m \mid x.$$

The resulting monitor LTS is called M^{alt} and it is a fragment of M^δ.

▶ **Corollary 22.** *If φ is a reliably monitorable formula on $(\textsc{Act}, \{\tau\})$, then there is an υ-alternating myopic monitor that monitors for φ on every LTS L on $(\textsc{Act}, \{\tau\})$ from every total obscuring of L.* ◀

▶ **Example 23.** We revisit the LTSs $L_1 \leq_o L_2$ from Example 19. Let $m_4 = \upsilon.\upsilon.\alpha.\mathtt{no}$ be a myopic monitor but *not* an υ-alternating one. We see that m_4 rejects process p_0 in L_1, but not in L_2 since m_4 flags processes that perform *at least two* silent actions before performing α; this is *not* the case for p_0 in L_2. The constraint of υ-alternation ensures that the monitors

are not allowed to count silent actions and thus rely on information that may be hidden in a further obscuring of the LTS: the information that a monitor of the form $\upsilon.\mathtt{no}$ or $\upsilon.\alpha.m$ can analyse is "at least one" silent transition (and then α in the second case), which is information guaranteed to be preserved by Properties E and D. ◀

The following results describe how monitor rejections are preserved by the obscuring preorder.

▶ **Lemma 24.** *For each $m \in \text{Mon}_\upsilon$ and each process p of L where $L \leq_o L'$, $\mathbf{rej}_{\langle M^\upsilon, I^\upsilon, L'\rangle}(m,p)$ implies $\mathbf{rej}_{\langle M^\upsilon, I^\upsilon, L\rangle}(m,p)$.* ◀

▶ **Lemma 25.** *For every υ-alternating myopic monitor m and p a process of L where $L \leq_o L'$, $\mathbf{rej}_{\langle M^\upsilon, I^\upsilon, L\rangle}(m,p)$ implies $\mathbf{rej}_{\langle M^\upsilon, I^\upsilon, L'\rangle}(m,p)$.* ◀

▶ **Corollary 26.** *If an υ-alternating myopic monitor monitors for a formula φ on LTS L on $(\text{Act}, \{\tau\})$ from any obscuring of L, then the monitor reliably monitors for φ on L.* ◀

We are ready to identify a maximal reliably monitorable fragment of μHML on $(\text{Act}, \{\tau\})$, which we call RsHML:

$$\theta, \chi \in \text{RsHML} ::= \mathtt{tt} \mid \mathtt{ff} \mid [\tau]\mathtt{ff} \mid [\alpha]\theta \mid [[\tau]][\alpha]\theta \mid \theta \wedge \chi \mid \max X.\theta \mid X.$$

▶ **Definition 27 (Reliable Monitor Synthesis).** We define a reliable monitor synthesis function $(\!|\cdot|\!)_r$ from RsHML to υ-alternating myopic monitors.

$$(\!|\mathtt{tt}|\!)_r = \mathtt{end} \qquad (\!|\mathtt{ff}|\!)_r = \mathtt{no} \qquad (\!|X|\!)_r = x \qquad (\!|[\tau]\mathtt{ff}|\!)_r = \upsilon.\mathtt{no}$$

$$(\!|\psi_1 \wedge \psi_2|\!)_r = \begin{cases} (\!|\psi_1|\!)_r & \text{if } (\!|\psi_2|\!)_r = \mathtt{end} \\ (\!|\psi_2|\!)_r & \text{if } (\!|\psi_1|\!)_r = \mathtt{end} \\ (\!|\psi_1|\!)_r + (\!|\psi_2|\!)_r & \text{otherwise} \end{cases} \qquad (\!|[\alpha]\psi|\!)_r = \begin{cases} \mathtt{end} & \text{if } (\!|\psi|\!)_r = \mathtt{end} \\ \alpha.(\!|\psi|\!)_r & \text{otherwise} \end{cases}$$

$$(\!|\max X.\psi|\!)_r = \begin{cases} \mathtt{end} & \text{if } (\!|\psi|\!)_r = \mathtt{end} \\ \mathtt{rec}\ x.(\!|\psi|\!)_r & \text{otherwise} \end{cases} \qquad (\!|[[\tau]][\alpha]\psi|\!)_r = \begin{cases} \mathtt{end} & \text{if } (\!|\psi|\!)_r = \mathtt{end} \\ \upsilon.\alpha.(\!|\psi|\!)_r & \text{otherwise} \end{cases}$$ ◀

▶ **Lemma 28.** *For every formula $\varphi \in \text{RsHML}$, $(\!|\varphi|\!)_r$ reliably monitors for φ.*

▶ **Definition 29 (Reliable Formula Synthesis).** We define a reliable formula synthesis function $\|\cdot\|_r$ from υ-alternating myopic monitors to RsHML.

$$\|\mathtt{end}\|_r = \mathtt{tt} \qquad \|\mathtt{no}\|_r = \mathtt{ff} \qquad \|x\|_r = X$$
$$\|\upsilon.\mathtt{no}\|_r = [\tau]\mathtt{ff} \qquad \|\upsilon.\alpha.m\|_r = [[\tau]][\alpha]\|m\|_r \qquad \|\alpha.m\|_r = [\alpha]\|m\|_r$$
$$\|m + n\|_r = \|m\|_r \wedge \|n\|_r \qquad \|\mathtt{rec}\ x.m\|_r = \max X.\|m\|_r$$ ◀

▶ **Lemma 30.** *For every υ-alternating myopic monitor m on $(\text{Act}, \{\tau, \upsilon\})$, m reliably monitors for $\|m\|_r$.* ◀

Theorem 31 presents the main result of this section. The first part follows from Lemmata 28 and 30 whereas the second part is a consequence of Lemma 30 and Corollaries 22 and 26.

▶ **Theorem 31.** *The monitoring system (M^{alt}, I^υ) on $(\text{Act}, \{\tau, \upsilon\})$ reliably monitors for* RsHML *on $(\text{Act}, \{\tau, \upsilon\})$. Moreover,* RsHML *is the largest reliably monitorable fragment of μHML up to logical equivalence.* ◀

We note that RsHML is also a fragment of ssHML, the maximally monitorable fragment of μHML identified in Section 3. Theorem 31 holds for every preorder $\leq_o$ that has the properties listed in Subsection 4.2. Therefore, it also holds for $\leq_c$ from Definition 16.

▶ **Example 32.** We return to the server from Examples 12 and 17. Notice that formula ψ is a RsHML-formula, and therefore it is reliably monitorable. ◀

We conclude by noting that myopic monitors can also be described as a fragment of full monitors by replacing $\upsilon.\texttt{no}$ by $\tau.\texttt{no}+\upsilon.no$ and $\upsilon.\alpha.m$ with $\texttt{rec}\ x.(\tau.x+\upsilon.x+\tau.\alpha.m+\upsilon.\alpha.m)$ in any myopic monitor. However, myopic monitors provide a cleaner, more efficient description of the same monitors.

6 Conclusions

We developed a general framework for *reliable monitorability* for monitoring setups that obscure information about the internal behaviour of systems. The framework is described through a family of LTS preorders that satisfy natural properties (A through D of Subsection 4.2). Via further assumptions (properties F, G, H) that guarantee a certain level of information obscuring, we identified RsHML, a maximal *reliably* monitorable fragment of μHML. Then, we provided a monitoring system, (M^{alt}, I^{υ}), that reliably monitors for RsHML.

Related work: In [9], Dwyer et al. use the approach of combining several properties to be monitored to produce a composite property and then project this composite property onto a smaller set of observable actions. This sampling technique effectively "silences" some of the observable actions and focuses on the rest to reduce overhead without risking unsound monitoring. Their approach highlights the importance of silent actions for RV and the need for a framework to handle imperfect information about silent events.

In [5] Basin *et. al.* consider the problem of monitoring over defective traces (called incomplete/disagreeing logs). They propose an augmented LTL specification language that permits reasoning about incompleteness and handling of inconsistencies. In some sense, this is related to our reliable monitors that are able to provide correct verdicts in the presence of event obscuring; however, the authors in [5] do not tackle issues related to monitorability.

In [19], Shi *et. al.* consider the problem of monitoring a wireless network via a wireless sniffer. A wireless sniffer may introduce uncertainty over a monitoring setup, as the trace it detects may not necessarily be the actual trace of the system, due to the intrinsic unreliability of the wireless network. The authors in [19] thus develop a monitoring framework to tolerate such errors. In separate work [6], Basin *et. al.* tackle the problem of distributed monitoring over a network which may produce delays and/or failures; they use a monitoring system based on a real-time three-valued logic that can track when an event took place. Their monitors may then need to draw appropriate conclusions under incomplete or scrambled information. Although we do not consider aspects such as event reordering, our work could serve as a basis for a better understanding of the level of obscuring these systems can tolerate.

Variations: Our system can be adjusted to describe diverse situations, by either weakening or strengthening the power of obscuring preorders. For instance, one can relax properties F to H to describe situations where there is a guarantee that the system reveals its internal behaviour at least partly and to a certain degree. For example, we can take $\leq_o$ to be the identity relation on LTSs, as it satisfies properties A to D and therefore is an obscuring preorder; then, the reliably monitorable properties would be all of ssHML and the corresponding monitoring system would be that of full monitors. A preorder between $\leq_c$ and $=$ would perhaps be more interesting.

References

1 Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfsdóttir, and Sævar Örn Kjartansson. Determinizing monitors for HML with recursion. *arXiv preprint arXiv:1611.10212*, 2016.

2 Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfsdóttir, and Sævar Örn Kjartansson. On the complexity of determinizing monitors. In Arnaud Carayol and Cyril Nicaud, editors, *Implementation and Application of Automata: 22nd International Conference, CIAA 2017, Marne-la-Vallée, France, June 27-30, 2017, Proceedings*, pages 1–13, Cham, 2017. Springer International Publishing. URL: http://dx.doi.org/10.1007/978-3-319-60134-2_1, doi:10.1007/978-3-319-60134-2_1.

3 Luca Aceto, Anna Ingólfsdóttir, Kim Guldstrand Larsen, and Jiri Srba. *Reactive Systems: Modelling, Specification and Verification*. Cambridge University Press, New York, NY, USA, 2007.

4 S. Arun-Kumar and Matthew Hennessy. An efficiency preorder for processes. *Acta Informatica*, 29(8):737–760, 1992. URL: http://dx.doi.org/10.1007/BF01191894, doi:10.1007/BF01191894.

5 David Basin, Felix Klaedtke, Srdjan Marinovic, and Eugen Zălinescu. Monitoring compliance policies over incomplete and disagreeing logs. In *Runtime Verification*, pages 151–167, 2013. URL: http://dx.doi.org/10.1007/978-3-642-35632-2_17, doi:10.1007/978-3-642-35632-2_17.

6 David Basin, Felix Klaedtke, and Eugen Zalinescu. Failure-aware Runtime Verification of Distributed Systems. In *35th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2015)*, pages 590–603, 2015. URL: http://drops.dagstuhl.de/opus/volltexte/2015/5619, doi:http://dx.doi.org/10.4230/LIPIcs.FSTTCS.2015.590.

7 Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. *Theoretical Computer Science*, 669:33–58, 2017.

8 Edmund M. Clarke, Jr., Orna Grumberg, and Doron A. Peled. *Model Checking*. MIT Press, Cambridge, MA, USA, 1999.

9 M. B. Dwyer, M. Diep, and S. Elbaum. Reducing the cost of path property monitoring through sampling. In *Proceedings of the 2008 23rd IEEE/ACM International Conference on Automated Software Engineering*, pages 228–237, 2008. URL: http://dx.doi.org/10.1109/ASE.2008.33, doi:10.1109/ASE.2008.33.

10 Adrian Francalanza. A Theory of Monitors. In *Foundations of Software Science and Computation Structures*, pages 145–161, 2016.

11 Adrian Francalanza, Luca Aceto, and Anna Ingolfsdottir. On verifying Hennessy-Milner logic with recursion at runtime. In Ezio Bartocci and Rupak Majumdar, editors, *Runtime Verification*, volume 9333 of *Lecture Notes in Computer Science*, pages 71–86. Springer International Publishing, 2015. URL: http://dx.doi.org/10.1007/978-3-319-23820-3_5, doi:10.1007/978-3-319-23820-3_5.

12 Adrian Francalanza, Luca Aceto, and Anna Ingolfsdottir. Monitorability for the Hennessy–Milner logic with recursion. *Formal Methods in System Design*, pages 1–30, 2017.

13 Adrian Francalanza and Aldrin Seychell. Synthesising Correct Concurrent Runtime Monitors. *Formal Methods in System Design (FMSD)*, 46(3):226–261, 2015. URL: http://dx.doi.org/10.1007/s10703-014-0217-9, doi:10.1007/s10703-014-0217-9.

14 Matthew Hennessy. *Algebraic Theory of Processes*. MIT Press, 1988.

15 Dexter Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27(3):333–354, 1983. URL: http://www.sciencedirect.com/science/article/pii/0304397582901256, doi:http://dx.doi.org/10.1016/0304-3975(82)90125-6.

16 Kim Guldstrand Larsen. Proof systems for satisfiability in Hennessy-Milner logic with recursion. *Theoretical Computer Science*, 72(2&3):265–288, 1990. URL: http://dx.doi.org/10.1016/0304-3975(90)90038-J, doi:10.1016/0304-3975(90)90038-J.

17 Martin Leucker and Christian Schallhart. A brief account of Runtime Verification. *The Journal of Logic and Algebraic Programming*, 78(5):293 – 303, 2009. URL: http://www.sciencedirect.com/science/article/pii/S1567832608000775, doi:http://dx.doi.org/10.1016/j.jlap.2008.08.004.

18 R. Milner. *Communication and concurrency*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1989.

19 Jinghao Shi, Shuvendu K. Lahiri, Ranveer Chandra, and Geoffrey Challen. Wireless protocol validation under uncertainty. In *Runtime Verification - 16th International Conference, RV 2016, Madrid, Spain, September 23-30, 2016, Proceedings*, pages 351–367, 2016. URL: http://dx.doi.org/10.1007/978-3-319-46982-9_22, doi:10.1007/978-3-319-46982-9_22.

20 Yoriyuki Yamagata, Cyrille Artho, Masami Hagiya, Jun Inoue, Lei Ma, Yoshinori Tanabe, and Mitsuharu Yamamoto. Runtime monitoring for concurrent systems. In *Runtime Verification - 16th International Conference, RV 2016, Madrid, Spain, September 23-30, 2016, Proceedings*, pages 386–403, 2016.

A Appendix

In this appendix, we prove the claims from the main body of this paper.

A.1 General Lemmata

We recursively define relation $\subseteq$ on monitors: for all $m, m' \in \text{Mon}_\delta$:

(i) $m \subseteq m$;

(ii) if $m \subseteq m'$, then
- $m \subseteq m' + n$ and
- $m \subseteq n + m'$; and

(iii) if $m \subseteq m'[\texttt{rec}\ x.m'/x]$, then $m \subseteq \texttt{rec}\ x.m'$.

Intuitively, if $m \subseteq m'$ and $m \xrightarrow{\mu} n$, then also $m' \xrightarrow{\mu} n$.

▶ **Lemma 33.** *If a $m \subseteq n$ for monitors m, n of monitor LTS M that occurs from the rules in Table 1, then $m \xrightarrow{\mu}_M m'$ implies $n \xrightarrow{\mu}_M m'$.*

Proof. By straightforward induction on the definition of $\subseteq$. ◀

▶ **Lemma 34.** *Let (M, I) be a monitoring system and R an instrumentation rule of the form*

$$R \frac{cond}{m \triangleleft p \xrightarrow{\mu}_{I(M,L)} \texttt{end} \triangleleft p'},$$

where cond is a sequence of conditions, m a monitor, p, p' processes from some LTS, and μ some action. Let N be the result of adding R to the instrumentation rules of I. Then, for every process p of an LTS L, and monitor m of M,

$$\mathbf{rej}_{\langle M,I,L\rangle}(m,p) \quad \textit{iff} \quad \mathbf{rej}_{\langle N,I,L\rangle}(m,p).$$

Proof. Notice that the effect of R on the LTS $I(N, L)$ is to add transitions leading to monitored processes of the form $\texttt{end} \triangleleft p$, which can never reach a monitored process of the form $\texttt{no} \triangleleft p'$. Therefore, a monitored process $m \triangleleft p$ can reach a monitored process of the form $\texttt{no} \triangleleft p'$ in $I(M, L)$ if and only if $m \triangleleft p$ can reach $\texttt{no} \triangleleft p'$ in $I(N, L)$. ◀

▶ **Lemma 35.** *Let (M, I^δ) be a monitoring system on (Act, Sil) and let m be a monitor of M. Then, for every process p of a system LTS L on (Act, Sil), $\mathbf{rej}_{\langle M,I^\delta,L\rangle}(m,p)$ if and only if there are some $s \in (\text{Act} \cup \text{Sil})^*$ and a process q of L, such that $p \stackrel{s}{\Rightarrow}_L q$ and $m \stackrel{s}{\Rightarrow}_M \texttt{no}$.*

Proof. The lemma is a consequence of the following claim: for every $s \in (\text{Act} \cup \text{Sil})^*$ and process q of L, $m \triangleleft p \stackrel{s}{\Rightarrow}_{I^\delta(M,L)} \texttt{no} \triangleleft q$ if and only if $p \stackrel{s}{\Rightarrow}_L q$ and $m \stackrel{s}{\Rightarrow}_M \texttt{no}$. The "if" direction can be proven by induction on $m \stackrel{s}{\Rightarrow}_M \texttt{no}$ and the "only if" direction by induction on $m \triangleleft p \stackrel{s}{\Rightarrow} \texttt{no} \triangleleft q$ (see also the proof of Proposition 1 in [11]). ◀

A.2 Omitted Proofs from Section 3

Theorem 11 was proven in [12] for the external monitoring system (M^α, I^α). As we demonstrate, it is not hard to recast that result for the full monitoring system and thus prove Theorem 8. For a formula $\varphi \in \text{ssHML}$, we define $w(\varphi)$ in the following way:

$$w(\texttt{tt}) = \texttt{tt}; \qquad w(\texttt{ff}) = \texttt{ff}; \qquad w(X) = X;$$
$$w(\varphi \wedge \psi) = w(\psi) \wedge w(\psi); \qquad w(\max X.\varphi) = \max X.w(\varphi);$$
$$w([\mu]\varphi) = [[\mu]]w(\varphi).$$

Proof of Theorem 8. For an LTS $L = \langle P, (\text{Act}, \text{Sil}), \rightarrow_L \rangle$, let $L^{\text{Sil}} = \langle P, (\text{Act} \cup \text{Sil}, \emptyset), \rightarrow_L \rangle$. For a monitor $m \in \text{Mon}_\delta$ and process $p \in P$,

$$\mathbf{rej}_{\langle M^\delta, I^\delta, L \rangle}(m, p) \text{ if and only if } \mathbf{rej}_{\langle M^\alpha, I^\alpha, L^{\text{Sil}} \rangle}(m, p),$$

as the same rules are used in both cases. Notice that for $\alpha \in \text{Act} \cup \text{Sil}$, $\stackrel{\alpha}{\Rightarrow}_{L^{\text{Sil}}} = \xrightarrow{\alpha}_{L^{\text{Sil}}} = \xrightarrow{\alpha}_L$, as there are no silent actions for L^{Sil}, so for every $\varphi \in \text{ssHML}$, $[\![\varphi]\!]_L = [\![w(\varphi)]\!]_{L^{\text{Sil}}}$.

Therefore, for every $\varphi \in \text{ssHML}$, we can define $(\!|\varphi|\!) = (\!|w(\varphi)|\!)_w$ and for every $m \in \text{Mon}_\delta$, $\|m\| = \|w^{-1}(m)\|_w$, where $(\!|\cdot|\!)_w$ and $\|\cdot\|_w$ are the monitor synthesis and formula synthesis functions for sHML as defined in [11]. Then, for every $\varphi \in \text{ssHML}$, $(\!|\varphi|\!)$ monitors for φ and, for every $m \in \text{Mon}_\delta$, m monitors for $\|m\|$. ◀

A.3 Omitted Proofs and Examples from Section 4

Proof of Proposition 18. Property F holds because moves 1 and 2 allow us to turn all τ-transitions to υ-transitions and move 3 allows us to make $\xrightarrow{\upsilon}$ transitive.

Property G holds by using move 1.

Property H is straightforward to verify: it is not hard to see how LTS L can be turned into L' from Example 13 using moves 3 to ensure that $p_i \xrightarrow{\upsilon} q_r$ and 4 to then remove $p_i \xrightarrow{\upsilon} q_1$.

Properties B through D are preservation properties and it suffices to prove them for $\leq_1$:

Properties B and C are straightforward to verify for each move.

Property A holds for $L \leq_1 L'$, as it is a general condition for every move.

For properties E and D, notice that moves 1, 2, and 3 preserve reachability with respect to $\xrightarrow{\tau} \cup \xrightarrow{\upsilon}$, so the properties hold; for move 4, if $p \xrightarrow{\tau}_L q$ (for either condition), then we are done; otherwise $p \xrightarrow{\upsilon}_L q$, so if $p \not\xrightarrow{\upsilon}_{L'} q$, then $q \not\xrightarrow{\alpha}$ for all $\alpha \in \text{Act}$ (so property E is verified) and there is some p' such that $p \xrightarrow{\upsilon}_{L'} p'$ and either $q \xrightarrow{\upsilon}_L p'$ or $q \xrightarrow{\tau}_L p'$, so property D is verified. ◀

▶ **Example 36.** To save notation, by $p \xrightarrow{\tau+\upsilon} q$ we mean $p \xrightarrow{\tau} q$ and $p \xrightarrow{\upsilon} q$; moreover, for $1 \leq i \leq 4$, by $L \leq_1^i L'$ we mean that $L \leq_1 L'$ by condition i of Definition 16. Recall LTSs L_1, L_2, and L_3 from Example 19; we also define the following LTSs:

$L_4 : p_0 \xrightarrow{\tau+\upsilon} p_1 \xrightarrow{\tau} p_2 \xrightarrow{\alpha} p_3$ $\quad$ $L_5 : p_0 \xrightarrow{\tau} p_1 \xrightarrow{\tau+\upsilon} p_2 \xrightarrow{\alpha} p_3$

$L'_4 : p_0 \xrightarrow{\upsilon} p_1 \xrightarrow{\tau} p_2 \xrightarrow{\alpha} p_3$ $\quad$ $L'_5 : p_0 \xrightarrow{\tau} p_1 \xrightarrow{\upsilon} p_2 \xrightarrow{\alpha} p_3$

$L^1_6 : p_0 \xrightarrow{\upsilon} p_1 \xrightarrow{\tau+\upsilon} p_2 \xrightarrow{\alpha} p_3$ $\quad$ $L^2_6 : p_0 \xrightarrow{\tau+\upsilon} p_1 \xrightarrow{\upsilon} p_2 \xrightarrow{\alpha} p_3$

$L^3_6 : p_0 \xrightarrow{\tau+\upsilon} p_1 \xrightarrow{\tau+\upsilon} p_2 \xrightarrow{\alpha} p_3$ $\quad$ $L_6 : p_0 \xrightarrow{\upsilon} p_1 \xrightarrow{\upsilon} p_2 \xrightarrow{\alpha} p_3$

$L_7 : p_0 \xrightarrow{\upsilon} p_1 \xrightarrow{\upsilon} p_2 \xrightarrow{\alpha} p_3$ and $p_0 \xrightarrow{\upsilon} p_2$ $\quad$ $L_8 : p_0 \xrightarrow{\tau+\upsilon} p_2 \xrightarrow{\alpha} p_3$ and $p_1 \xrightarrow{\tau} p_2$

$L'_8 : p_0 \xrightarrow{\upsilon} p_2 \xrightarrow{\alpha} p_3$ and $p_1 \xrightarrow{\tau} p_2$ $\quad$ $L_9 : p_0 \xrightarrow{\upsilon} p_2 \xrightarrow{\alpha} p_3$ and $p_1 \xrightarrow{\tau+\upsilon} p_2$

We observe that by using moves 1 and then 2:

$L_1 \leq_1^1 L_4$; $\quad L_1 \leq_1^1 L_5$; $\quad L_1 \leq_1^1 L_4$; $\quad L_4 \leq_1^1 L^3_6$; $\quad L_5 \leq_1^1 L^3_6$; $\quad L_3 \leq_1^1 L_8$;

$L_4 \leq_1^2 L'_4$; $\quad L_5 \leq_1^2 L'_5$; $\quad L_8 \leq_1^2 L'_8$; $\quad L_9 \leq_1^2 L_2$; $\quad L^i_6 \leq_1^2 L_6$,

where $i = 1, 2, 3$. Finally, by using moves 3 and 4, we can see that $L_6 \leq_1^3 L_7 \leq_1^4 L_2$. Therefore, $L_1 \leq_c L_2$ and $L_3 \leq_c L_2$, and we were able to construct L_2 as an obscuring of L_1 and L_3 (and of all L_i for $1 \leq i \leq 9$). ◀

A.4 Omitted Proofs and Examples from Section 5

We first examine reliably monitorable formulas of μHML *w.r.t.* their monitorability on total obscurings and identify a suitable fragment of full monitors that suffices to monitor for all reliably monitorable formulas on total obscurings (Lemma 41).

▶ **Lemma 37.** *If φ is reliably monitorable on LTSs L, L' on $(\text{Act}, \{\tau\})$ and for an LTS L'' on $(\text{Act}, \{\tau, \upsilon\})$, $L \leq_o L''$ and $L' \leq_o L''$, then $[\![\varphi]\!]_L = [\![\varphi]\!]_{L'}$.*

Proof. If m of monitoring system (M, I) reliably monitors for φ on L and L', $L \leq_o L''$, and $L' \leq_o L''$, then

$$[\![\varphi]\!]_L = [\![\varphi]\!]_{L'} = \{p \mid \text{not } \mathbf{rej}_{\langle M,I,L''\rangle}(m, p)\}.$$ ◀

Let $[\tau + \upsilon]\varphi$ be short for $[\tau]\varphi \wedge [\upsilon]\varphi$. For a formula φ of ssHML on $(\text{Act}, \{\tau\})$, φ^υ is the result of replacing all occurrences of $[\tau]$ by $[\tau + \upsilon]$ in φ.

▶ **Lemma 38.** *For each (possibly open) ssHML-formula ψ on $(\text{Act}, \{\upsilon\})$, LTS L on $(\text{Act}, \{\upsilon\})$, and environment ρ, $[\![\psi, \rho]\!]_{L^\tau} = [\![\psi^\upsilon, \rho]\!]_L$*

Proof. By straightforward induction on ψ. ◀

▶ **Lemma 39.** *If φ is a ssHML-formula on $(\text{Act}, \{\tau\})$ that is reliably monitorable, then for every LTS L on $(\text{Act}, \{\tau\})$ and every obscuring L' of L, $[\![\varphi]\!]_L = [\![\varphi^\upsilon]\!]_{L'}$.*

Proof. By property G, we have that $(L')^\tau \leq_o L'$. Therefore, by Lemma 37, $[\![\varphi]\!]_L = [\![\varphi]\!]_{(L')^\tau}$. Finally, by Lemma 38, $[\![\varphi]\!]_{(L')^\tau} = [\![\varphi^\upsilon]\!]_{L'}$. ◀

▶ **Corollary 40.** *If φ is a reliably monitorable formula on $(\text{Act}, \{\tau\})$, then there is a full monitor on $(\text{Act}, \{\upsilon\})$ that reliably monitors for φ.*

Proof. Since φ is reliably monitorable, then it is also monitorable by a monitor m. Using the formula synthesis function from Section 3, $\|m\| \equiv \varphi$ and $\|m\| \in$ ssHML (see the proof of Theorem 8 above). Thus, we assume that $\varphi \in$ ssHML. Then, the corollary is the result of Lemma 39. ◀

▶ **Lemma 41.** *For every reliably monitorable formula φ on $(\text{Act}, \{\tau\})$, there is an υ-alternating full monitor on $(\text{Act}, \{\tau, \upsilon\})$ that reliably monitors for φ on every LTS L on $(\text{Act}, \{\tau\})$ from any total obscuring of L.*

Proof. According to Corollary 40, there is a full monitor m on $(\text{Act}, \{\tau, \upsilon\})$ that reliably monitors for φ. We demonstrate that m rejects exactly the same processes as an υ-alternating full monitor on $(\text{Act}, \{\upsilon\})$ on total LTSs.

Let Q be the set of monitors that can be reached by a sequence of transitions from m. Let A be the NFA $(Q, \text{Act} \cup \{\upsilon, (\upsilon\alpha) \mid \alpha \in \text{Act}\}, \delta, m, \{\texttt{no}\})$,[1] where for $n \in Q$ and $\alpha \in \text{Act}$,

$$\begin{aligned} \delta(n, \alpha) &= \{n' \in Q \mid n \xrightarrow{\alpha} n'\}; \\ \delta(n, (\upsilon\alpha)) &= \{n' \in Q \mid n \xrightarrow{\upsilon}\xrightarrow{\alpha} n'\}; \text{ and} \\ \delta(n, \upsilon) &= \{\texttt{no} \mid n \xrightarrow{\upsilon} \texttt{no}\}. \end{aligned}$$

[1] For an automaton $(Q, \Sigma, \delta, q_0, F)$, Q is the set of states, Σ its alphabet, δ its transition relation, $q_0 \in Q$ its initial state, and $F \subseteq Q$ the set of its accepting states.

For a trace $s \in (\text{Act} \cup \{\upsilon\})^*$, we define $o'(s)$ as follows: $o'(\epsilon) = \epsilon$; $o'(\upsilon^k) = \upsilon$ for $k > 0$; $o'(\alpha s) = \alpha o'(s)$ for $\alpha \in \text{Act}$; $o'(\upsilon^k \alpha s) = (\upsilon\alpha)o'(s)$ for $k > 0$. Notice that $o(s)$ is $o'(s)$ without the parentheses.

Claim: for $s \in (\text{Act} \cup \{\upsilon\})^*$, $m \stackrel{s}{\Rightarrow} \texttt{no}$ if and only if A accepts $o'(s)$.

To see that the above claim holds, notice first of all that A accepts $o'(s)$ if and only if $m \stackrel{o(s)}{\Longrightarrow} \texttt{no}$. Therefore, to show the claim it suffices to prove that $m \stackrel{s}{\Rightarrow} \texttt{no}$ if and only if $m \stackrel{o(s)}{\Longrightarrow} \texttt{no}$. By property H of $\leq_o$, there is a process p of some LTSs $L \leq_o L'$, such that p produces exactly the prefixes of s in L and the prefixes of $o(s)$ in L'. Monitor m monitors for φ on $L^\tau \leq_o L$ (by property G) from both L and L' and therefore, $m \stackrel{s}{\Rightarrow} \texttt{no}$ if and only if $m \stackrel{o(s)}{\Longrightarrow} \texttt{no}$.

By [1, Theorem 6], there is an external monitor n' on $(\text{Act} \cup \{\upsilon, (\upsilon\alpha) \mid \alpha \in \text{Act}\}, \emptyset)$ which rejects exactly the traces A accepts; to turn this into an equivalent (in that it rejects $o(s)$ iff n' rejects $o'(s)$) υ-alternating full monitor n on (Act, υ), we simply replace all $(\upsilon\alpha)$ by $\upsilon.\alpha$ in n'. Then, for every trace $s \in (\text{Act} \cup \{\upsilon\})^*$, $m \stackrel{s}{\Rightarrow} \texttt{no}$ if and only if $n \stackrel{o(s)}{\Longrightarrow} \texttt{no}$.

Now we are ready to prove that m and n reject the same processes on all total LTSs L. If m rejects process p, then, by Lemma 35, p can produce a trace s, which m rejects; but $\stackrel{\upsilon}{\rightarrow}$ is transitive in L and therefore, p can also produce trace $o(s)$, which n rejects. The other direction is similar. ◀

Now we know that any monitoring system that reliably monitors for a fragment of μHML on $(\text{Act}, \{\tau\})$ must be equivalent to a fragment of full monitors on $(\text{Act}, \{\upsilon\})$. We can further narrow down the kind of monitoring system we are looking for by observing that we can syntactically restrict full monitors and still preserve monitorability with respect to reliably monitorable formulas.

▶ **Definition 42.** The υ-alternating full monitors on $(\text{Act}, \{\upsilon\})$ are the full monitors described by the following grammar:

$$m, n \in \text{Mon}_{alt} ::= \texttt{end} \mid \texttt{no} \mid \upsilon.\texttt{no} \mid \alpha.m \mid \upsilon.\alpha.m \mid m{+}n \mid \texttt{rec}\ x.m \mid x$$

The resulting monitor LTS is called M^{alt} and it is a fragment of M^δ. ◀

▶ **Example 43.** We revisit Example 19. Observe that $m_3^\upsilon = \upsilon.\alpha.\texttt{no}$ rejects the same processes as m_3 in L_2 (and every other total obscuring). m_3^υ is also an υ-alternating full monitor and it monitors for φ_3 on total obscurings of LTSs, but it fails to monitor for φ_3 on L_1 (even after replacing all τ-transitions by υ-transitions). Thus, Lemma 41 does not identify a class of monitors that reliably monitor for formulas, as even υ-alternating full monitors are too powerful: m_3^υ monitors for the property that a process cannot perform *one* silent action followed by an α, so it still counts silent actions. Lemma 41, however, serves as a guide for identifying such a class, as the following subsection demonstrates. ◀

We define υ_o-monitors as a technical tool. Just like myopic monitors, υ_o-monitors use the syntax of full monitors on $(\text{Act}, \{\upsilon\})$ and thus

$$M^{\upsilon_o} = \langle \text{Mon}_\upsilon, (\text{Act}, \text{Sil}), \rightarrow \rangle.$$

For the instrumentation I^{υ_o} of υ_o-monitors, we use the same rules as for myopic monitors, from Table 2, except for iTran.

For an LTS L on $(\text{Act}, \{\tau, \upsilon\})$, L^υ is the LTS with the same set of processes and actions, but with the following change in the transition relation: for $\delta \in \text{Act} \cup \{\tau\}$, $\stackrel{\delta}{\rightarrow}_{L^\upsilon} = \stackrel{\delta}{\rightarrow}_L$; and $\stackrel{\upsilon}{\rightarrow}_{L^\upsilon} = \stackrel{\upsilon}{\rightarrow}_L \cup \stackrel{\tau}{\rightarrow}_L$.

▶ **Lemma 44.** *For every monitor $m \in \text{Mon}_\upsilon$, trace $s \in (\text{Act} \cup \{\tau, \upsilon\})^*$, and all processes p, q of an LTS L on $(\text{Act}, \{\tau, \upsilon\})$,*

$$m \triangleleft p \stackrel{s}{\Rightarrow}_{I^{\upsilon_o}(M^{\upsilon_o}, L)} \mathtt{no} \triangleleft q \quad \textit{iff} \quad m \triangleleft p \stackrel{s}{\Rightarrow}_{I^{\delta}(M^{\delta}, L^{\upsilon})} \mathtt{no} \triangleleft q.$$

Proof. By Lemma 34, we can examine the monitoring systems without rule iTer. Observe that in L^υ, rule iMonS is a special case of iMon, because $\stackrel{\tau}{\to}_{L^\upsilon} \subseteq \stackrel{\upsilon}{\to}_{L^\upsilon}$. Therefore, the instrumentations of the two monitoring systems give exactly the same instrumentation LTSs. ◀

▶ **Corollary 45.** *For every monitor m and LTS L on $(\text{Act}, \{\tau, \upsilon\})$,*

$$\mathbf{rej}_{\langle M^{\upsilon_o}, I^{\upsilon_o}, L\rangle}(m, p) \quad \textit{iff} \quad \mathbf{rej}_{\langle M^{\delta}, I^{\delta}, L^{\upsilon}\rangle}(m, p).$$

▶ **Lemma 46.** *If in an LTS L on $(\text{Act}, \{\tau, \upsilon\})$, $\stackrel{\tau}{\to}_L \subseteq \stackrel{\upsilon}{\to}_L$ and $\stackrel{\upsilon}{\to}_L$ is transitive, then for every $m \in \text{Mon}_\upsilon$ and process p of L,*

$$\mathbf{rej}_{\langle M^{\upsilon_o}, I^{\upsilon_o}, L\rangle}(m, p) \quad \textit{iff} \quad \mathbf{rej}_{\langle M^{\upsilon}, I^{\upsilon}, L\rangle}(m, p).$$

Proof. Notice that in L, rule iTran is redundant: if $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L)} n \triangleleft q$ is constructed after an application of rule iMon and k applications of rule iTran, then $m \stackrel{\upsilon}{\to}_{M^\upsilon} n$ and $p(\stackrel{\upsilon}{\to}_L \cup \stackrel{\tau}{\to}_L)^{k+1} q$. But $\stackrel{\tau}{\to}_L = \emptyset$ and $\stackrel{\upsilon}{\to}_L$ is transitive, so $p \stackrel{\upsilon}{\to}_L q$ and therefore, $m \triangleleft p \stackrel{\upsilon}{\to}_{I^{\upsilon_o}(M^\upsilon, L)} n \triangleleft q$ (and $M^{\upsilon_o} = M^\upsilon$ by definition). ◀

Proof of Lemma 21. By property F of $\leq_o$, notice that if L is total, then $L^\upsilon = L$, $\stackrel{\tau}{\to}_L$ is empty, and $\stackrel{\upsilon}{\to}_L$ is transitive. The corollary then is a consequence of Corollary 45 and Lemma 46. ◀

Proof of Corollary 22. A direct consequence of Lemmata 41 and 21. ◀

Proof of Lemma 24. To prove the lemma, it suffices to prove that if $m \triangleleft p \stackrel{\delta}{\to}_{I^\upsilon(M^\upsilon, L')} m' \triangleleft p'$, then $m \triangleleft p \stackrel{\delta}{\to}_{I^\upsilon(M^\upsilon, L)} m' \triangleleft p'$, where $\delta \in \text{Act} \cup \{\tau, \upsilon\}$. By Lemma 34, we can ignore transitions derived from rule iTer. By properties A and B of $\leq_o$, for $\delta \in \text{Act} \cup \{\tau\}$, $\stackrel{\delta}{\to}_{L'} \subseteq \stackrel{\delta}{\to}_L$, so $\stackrel{\delta}{\to}_{I^\upsilon(M^\upsilon, L')} \subseteq \stackrel{\delta}{\to}_{I^\upsilon(M^\upsilon, L)}$ — as besides rule iTer, that we ignore, the conditions of all other instrumentation rules are positive on $\stackrel{\tau}{\to}$.

It remains to prove the claim that if $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L')} m' \triangleleft p'$, then $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L)} m' \triangleleft p'$. Transition $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L')} m' \triangleleft p'$ may be the result of applying rule iMon to derive an υ-transition and then use zero or more applications of rule iTran. We prove the claim by induction on the number of applications of rule iTran.

Base Case: Rule iMon was applied on $m \triangleleft p$, so $m \stackrel{\upsilon}{\to}_{M^\upsilon} m'$ and $p \stackrel{\mu}{\to}_{L'} p'$ for some $\mu \leq \upsilon$ — that is, μ is either τ or υ. If $\mu = \tau$, then by property B, $p \stackrel{\mu}{\to}_L p'$ and therefore we can use rule iMon to produce $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L)} m' \triangleleft p'$. If $\mu = \upsilon$, then $p \stackrel{\upsilon}{\to}_{L'} p'$ and $m \stackrel{\upsilon}{\to}_{M^\upsilon} m'$. By property C, $p \stackrel{\mu'}{\to}_L p''(\stackrel{\upsilon}{\to}_L \cup \stackrel{\tau}{\to}_L)^* p'$ for some process p'' and $\mu' \in \{\upsilon, \tau\}$. Therefore, we can use rule iMon to derive $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L)} m' \triangleleft p''$ and rule iTran 0 or more times to derive $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L)} m' \triangleleft p'$.

Inductive Step: If rule iTran was applied $k > 0$ times to derive transition $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L')} m' \triangleleft p'$, then it was used $k-1$ times to derive $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L')} m' \triangleleft p''$, for some $p'' \stackrel{\mu}{\to}_{L'} p'$, where $\mu \in \{\tau, \upsilon\}$. By the inductive hypothesis, it is the case that $m \triangleleft p \stackrel{\upsilon}{\to}_{I^\upsilon(M^\upsilon, L)} m' \triangleleft p''$. If $p'' \stackrel{\tau}{\to}_{L'} p'$, then by property B, also $p'' \stackrel{\tau}{\to}_L p'$, so we can use rule iTran and derive

$m \triangleleft p \xrightarrow{\upsilon}_{I^\upsilon(M^\upsilon,L')} m' \triangleleft p'$. On the other hand, if $p'' \xrightarrow{\upsilon}_{L'} p'$, then by property C, $p''(\xrightarrow{\tau}_L \cup \xrightarrow{\upsilon}_L)^+ p'$ and we can again use rule iTran, perhaps more than once, and derive $m \triangleleft p \xrightarrow{\upsilon}_{I^\upsilon(M^\upsilon,L')} m' \triangleleft p'$. ◀

▶ **Lemma 47.** *For all m, m', υ-alternating myopic monitors, $m \xrightarrow{\upsilon}_p m'$ if and only if either $m' = \mathtt{no}$ and $\mathtt{no} \subseteq m$ or $\upsilon.m' \subseteq m$.*

Proof. The "if" direction is a consequence of Lemma 33 and rule mVrd. For the "only if" direction, notice that $m \xrightarrow{\upsilon}_{M^\upsilon} m'$ can only be produced by an application of rule mAct or mVrd and k applications of a combination of rules mSel and mRec. The remaining proof is by straightforward induction on k. ◀

Proof of Lemma 25. The proof is by induction on the number of visible actions on a rejecting path of the instrumentation. If $m = \mathtt{no}$, then we are done, so in the following we assume that $m \neq \mathtt{no}$. If $m \triangleleft p(\xrightarrow{\upsilon}_{I^\upsilon(M^\upsilon,L)})^+ \mathtt{no} \triangleleft p'$, then $p(\xrightarrow{\tau}_L \cup \xrightarrow{\upsilon}_L)^+ p'$ and $m \xrightarrow{\upsilon}_{M^\upsilon} m'(\xrightarrow{\upsilon}_{M^\upsilon})^* \mathtt{no}$; therefore, $p \xrightarrow{\tau}_L \cup \xrightarrow{\upsilon}_L$ and by Lemma 47, $\upsilon.m' \subseteq m$, so $m' = \mathtt{no}$ by the syntactic restrictions of υ-alternating myopic monitors. By property D, $p \xrightarrow{\tau}_{L^\upsilon} \cup \xrightarrow{\upsilon}_{L^\upsilon} p''$ for some p'' and therefore, $m \triangleleft p \xrightarrow{\upsilon}_{I^\upsilon(M^\upsilon,L)} \mathtt{no} \triangleleft p''$. On the other hand, if

$$m \triangleleft p(\xrightarrow{\upsilon}_{I^\upsilon(M^\upsilon,L)})^+ n \triangleleft q \xrightarrow{\alpha}_{I^\upsilon(M^\upsilon,L)} n' \triangleleft q' \overset{s}{\Rightarrow}_{I^\upsilon(M^\upsilon,L)} m' \triangleleft p',$$

where $\alpha \in \text{Act}$, then either $\upsilon.\mathtt{no} \subseteq m$ and we can complete the proof as above, or $\upsilon.\alpha.n' \subseteq m$. In the second case, $p(\xrightarrow{\tau}_L \cup \xrightarrow{\upsilon}_L)^+ q \xrightarrow{\alpha}_L q'$ and by the inductive hypothesis, $\mathbf{rej}_{\langle M^\upsilon, I^\upsilon, L'\rangle}(n', q')$. By property A, $q \xrightarrow{\alpha}_{L'} q'$ and by property E, $p(\xrightarrow{\tau}_{L'} \cup \xrightarrow{\upsilon}_{L'})^+ q$. Therefore, using rules iMon and iTran, $m \triangleleft p \xrightarrow{\upsilon}_{I^\upsilon(M^\upsilon,L')} \alpha.n' \triangleleft q \xrightarrow{\alpha}_{I^\upsilon(M^\upsilon,L')} n' \triangleleft q'$, and thus $\mathbf{rej}_{\langle M^\upsilon, I^\upsilon, L'\rangle}(m, p)$. ◀

Proof of Lemma 28. Because of Corollary 26, it suffices to prove that $(\!|\varphi|\!)_r$ monitors for φ on every LTS L on (Act, τ). Let L^+ be the LTS on $(\text{Act}, \{\tau, \upsilon\})$ which has the same processes as L, such that $\xrightarrow{\delta}_{L^+} = \xrightarrow{\delta}_L$ for $\delta \in \text{Act} \cup \{\tau\}$ and $\xrightarrow{\upsilon}_{L^+} = \overset{\tau}{\Rightarrow}_L$. We can immediately see that $\overset{\tau}{\Rightarrow}_{L^+} = \overset{\tau}{\Rightarrow}_L$ and therefore $[\![\varphi]\!]_L = [\![\varphi]\!]_{L^+}$. Let φ^s be the formula of ssHML on $(\text{Act} \cup \{\upsilon\}, \emptyset)$ which results from replacing $[\tau].\mathtt{no}$ by $[\upsilon].\mathtt{no}$ and $[[\tau]]$ by $[\upsilon]$. Then, $(\!|\varphi|\!)_r = (\!|\varphi^s|\!)$ and $[\![\varphi]\!]_{L^+} = [\![\varphi^s]\!]_{L^+}$ (by straightforward induction on φ); therefore, $(\!|\varphi|\!)_r$ (M^δ, I^δ)-monitors for φ on L^+ and by Corollary 45 and Lemma 46, $(\!|\varphi|\!)_r$ (M^υ, I^υ)-monitors for φ on L^+.

It now suffices to prove that $(\!|\varphi|\!)_r$ rejects a process in L if and only if it rejects the process in L^+. If $(\!|\varphi|\!)_r$ rejects a process p in L, then since p can produce the same trace that $(\!|\varphi|\!)_r$ rejects (Lemma 35) in L^+ as well, so $(\!|\varphi|\!)_r$ also rejects p in L^+. For the other direction, we use induction on k, where by Lemma 46,

$$(\!|\varphi|\!)_r \triangleleft p \xrightarrow{\mu_1}_{I^{\upsilon_o}(M^{\upsilon_o},L^+)} \cdots \xrightarrow{\mu_k}_{I^{\upsilon_o}(M^{\upsilon_o},L^+)} \mathtt{no} \triangleleft p' :$$

if $k = 0$, then we are done; if $k > 0$ and $\mu_1 \neq \upsilon$, then we are done by the inductive hypothesis; otherwise, by the inductive hypothesis, $(\!|\varphi|\!)_r \triangleleft p \xrightarrow{\upsilon}_{I^{\upsilon_o}(M^{\upsilon_o},L^+)} m \triangleleft q \overset{s}{\Rightarrow}_{I^\upsilon(M^\upsilon,L)} \mathtt{no} \triangleleft p'$ for some m, q, s. But $(\!|\varphi|\!)_r \triangleleft p \xrightarrow{\upsilon}_{I^{\upsilon_o}(M^{\upsilon_o},L^+)} m \triangleleft q$ implies that $(\!|\varphi|\!)_r \xrightarrow{\upsilon}_{M^{\upsilon_o}} m$ and $p \xrightarrow{\upsilon}_{L^+} q$; therefore, $(\!|\varphi|\!)_r \xrightarrow{\upsilon}_{M^\upsilon} m$ and $p(\xrightarrow{\tau}_L)^+ q$ and by using rule iMon and rule iTran 0 or more times, $(\!|\varphi|\!)_r \triangleleft p \xrightarrow{\upsilon}_{I^\upsilon(M^\upsilon,L)} m \triangleleft q$ and the induction is complete. ◀

Proof of Lemma 30. Similar to the proof of Lemma 28. ◀

Proof of Theorem 31. That (M^{alt}, I^υ) reliably monitors for RsHML is a direct consequence of Lemmata 28 and 30. We demonstrate that if φ is a reliably monitorable formula, then

it is equivalent to a RsHML formula. If φ is reliably monitorable, then by Corollary 22, there is an v-alternating myopic monitor m that monitors for φ on every LTS from every total obscuring of L. By Corollary 26, for every LTS L, m reliably monitors for φ from L — therefore, m reliably monitors for φ. By Lemma 30, m also reliably monitors for $\|m\|_r \in$ RsHML and therefore, for every LTS L,

$$[\![\varphi]\!]_L \;=\; \{p \mid \text{ not } \mathbf{rej}_{\langle M^{alt}, I^v, L\rangle}(m,p)\} \;=\; [\![\|m\|_r]\!]_L. \qquad \blacktriangleleft$$