
BIDIRECTIONAL RUNTIME ENFORCEMENT OF FIRST-ORDER BRANCHING-TIME PROPERTIES

LUCA ACETO, IAN CASSAR, ADRIAN FRANCALANZA, AND ANNA INGÓLFSDÓTTIR

^a ICE-TCS, Department of Computer Science, Reykjavík University, Iceland and Gran Sasso Science Institute, L'Aquila, Italy
e-mail address: luca@ru.is

^b Department of Computer Science, University of Malta, Malta and ICE-TCS, Department of Computer Science, Reykjavík University, Iceland
e-mail address: ian.cassar.10@um.edu.mt

^c Department of Computer Science, University of Malta, Malta
e-mail address: adrian.francalanza@um.edu.mt

^d ICE-TCS, Department of Computer Science, Reykjavík University, Iceland
e-mail address: anna@ru.is

ABSTRACT. Runtime enforcement is a dynamic analysis technique that instruments a monitor with a system in order to ensure its correctness as specified by some property. This paper explores *bidirectional* enforcement strategies for properties describing the input and output behaviour of a system. We develop an operational framework for bidirectional enforcement and use it to study the enforceability of the safety fragment of Hennessy-Milner logic with recursion (sHML). We provide an automated synthesis function that generates correct monitors from sHML formulas, and show that this logic is enforceable via a specific type of bidirectional enforcement monitors called action disabling monitors.

Key words and phrases: runtime monitors, property enforcement, monitor synthesis, first-order safety properties, modal μ -calculus.

“TheoFoMon: Theoretical Foundations for Monitorability” (nr.163406-051) of the Icelandic Research Fund and the Italian MIUR project PRIN 2017FTXR7S IT MATTERS “Methods and Tools for Trustworthy Smart Systems”.

“Developing Theoretical Foundations for Runtime Enforcement” (nr.184776-051) of the Icelandic Research Fund and the Endeavour Scholarship Scheme (Malta), part-financed by the European Social Fund (ESF) - Operational Programme II – 2014-2020.

“TheoFoMon: Theoretical Foundations for Monitorability” (nr.163406-051) and “MoVeMnt: Mode(1)s of Verification and Monitorability” (nr.217987-051) of the Icelandic Research Fund and the EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement nr. 778233.

“TheoFoMon: Theoretical Foundations for Monitorability” (nr.163406-051) of the Icelandic Research Fund.

1. INTRODUCTION

Runtime enforcement (RE) [LBW05, FFM08] is a dynamic verification technique that uses *monitors* to analyse the runtime behaviour of a system-under-scrutiny (SuS) and transform it in order to conform to some correctness *specification*. The earliest work in RE [LBW05, LBW09, Sak09, BM11a, KT12] models the behaviour of the SuS as a *trace* of *abstract* actions (e.g., $\alpha, \beta, \dots \in \text{ACT}$). Importantly, it assumes that the monitor can either suppress or replace *any* of these (abstract) actions and, whenever possible, insert additional actions into the trace.

This work has been used as a basis to implement *unidirectional* enforcement approaches [KAB⁺17, FFM12, ACFI18, Av11] that monitor the outputted trace of actions emitted by the SuS as illustrated by Figure 1(a). In this setup, the monitor is instrumented with the SuS to form a *composite system* (represented by the dashed enclosure in Figure 1) and is tasked with transforming the output behaviour of the SuS to ensure its correctness. For instance, an erroneous output β of the SuS is intercepted by the monitor and transformed into β' , to stop the error from propagating to the surrounding environment.

Despite its merits, unidirectional enforcement disregards the fact that not all events originate from the SuS. For instance, protocol specifications describing the interaction of communicating computational entities include *input* actions, instigated by the environment in addition to output actions originating from the SuS. Arguably, these properties are harder to enforce. Since inputs are instigated by the environment, the SuS possesses only partial control over them and the capabilities to prevent or divert such actions can be curtailed. Moreover, in such a bidirectional setting, the properties to be enforced tend to be of a *first-order* nature [ACFI18, HRTZ18], describing relationships between the respective payload carried by input and output events. This means that even when the (monitored) SuS can control when certain inputs can be supplied (e.g., by opening a communication port, or by reading a record from a database *etc.*), the environment still has control over the provided payload.

Broadly, there are two approaches to enforce bidirectional properties at runtime. Several bodies of work employ two monitors attached at the output side of each (diadic) interacting party [BCD⁺17, JGP16, CBD⁺12, ?]. As shown in Figure 1(b), the extra monitor is attached to the *environment* to analyse its outputs before they are passed on as *inputs* to the SuS. While this approach is effective, it assumes that a monitor can actually be attached to the environment (which is often inaccessible).

By contrast, Figure 1(c) presents a less explored *bidirectional enforcement* approach where the monitor analyses the entire behaviour of the SuS without the need to instrument the environment. The main downside of this alternative setup is that it enjoys limited control over the SuS’s inputs. As we already argued, the monitor may be unable to enforce a property that could be violated by an input action with an invalid payload value. In other cases, the monitor might need to adopt a different enforcement strategy to the ones that are conventionally used for enforcing output behaviour in a unidirectional one.

This paper explores how existing monitor transformations—namely, suppressions, insertions and replacements—can be repurposed to work for bidirectional enforcement, *i.e.*, the setup in Figure 1(c). Since inputs and outputs must be enforced differently, we find it essential to distinguish between the monitor’s transformations and their resulting effect on the visible behaviour of the composite system. This permits us to study the enforceability of properties defined via a safety subset of the well-studied branching-time logic

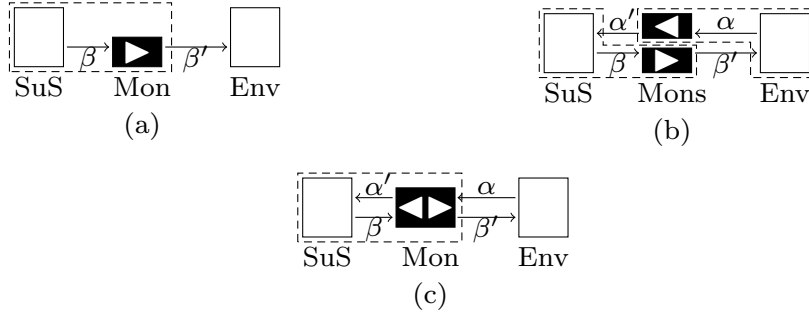


FIGURE 1. Enforcement instrumentation setups.

μ HML [RH97, AILS07, Lar90] (a reformulation of the modal μ -calculus [Koz83]), called SHML. A crucial aspect of our investigation is the *synthesis* function that maps the *declarative* safety μ HML specifications to *algorithmic* monitors that enforce properties concerning both the input and output behaviour of the SuS. Since monitors are part of the trusted computing base, it was imperative that we ensure that all synthesised monitors are correct [?, Fra21]. Our contributions are:

- (i) A general instrumentation framework for bidirectional enforcement (Figure 4) that is parametrisable by any system whose behaviour can be modelled as a labelled transition system. The framework subsumes the one presented in previous work [ACFI18] and differentiates between the enforcement of input and output actions.
- (ii) A novel formalisation describing what it means for a monitor to *adequately enforce* a property in a bidirectional setting (Definitions 4.1 and 4.9). These definitions are parametrisable with respect to an instrumentation relation, an instance of which is given by our enforcement framework of Figure 4.
- (iii) A new result showing that the subclass of *disabling monitors*, Definition 3.1 (the counterpart to suppression monitors in unidirectional enforcement), suffices to bidirectionally enforce safety properties expressed as μ HML formulas (Theorem 5.5). A by-product of this result is a synthesis function (Definition 5.3) that generates a disabling monitor from such safety formulas.
- (iv) A preliminary investigation on the notion of monitor *optimality* (Definition 6.3). Our proposed definition assesses the level of intrusiveness of the monitor and guides in the search for the least intrusive one. We evaluate our monitor synthesis function of Definition 5.3 in terms of this optimality measure, Theorem 6.12.

This article is the extended version of the paper titled “*On Bidirectional Runtime Enforcement*” that appeared at FORTE 2021 [ACFI21]. In addition to the material presented in the conference version, this version contains extended examples, the proofs of the main results and new material on monitor optimality. The related work section has also been considerably expanded.

2. PRELIMINARIES

The Model. We assume a countably infinite set of communication ports $a, b, c \in \text{PORT}$, a set of values $v, w \in \text{VAL}$, and partition the set of actions ACT into

Syntax			
$\varphi, \psi \in \text{sHML} ::= \text{tt}$	(truth)	ff	(falsehood)
$[p, c]\varphi$	(necessity)	$\max X.\varphi$	(greatest fp.)
		$\bigwedge_{i \in I} \varphi_i$	(conjunction)
		X	(fp. variable)
Semantics			
$\llbracket \text{tt}, \rho \rrbracket \stackrel{\text{def}}{=} \text{SYS}$			
$\llbracket \text{ff}, \rho \rrbracket \stackrel{\text{def}}{=} \emptyset$			
$\llbracket X, \rho \rrbracket \stackrel{\text{def}}{=} \rho(X)$			
$\llbracket \bigwedge_{i \in I} \varphi_i, \rho \rrbracket \stackrel{\text{def}}{=} \bigcap_{i \in I} \llbracket \varphi_i, \rho \rrbracket$			
$\llbracket \max X.\varphi, \rho \rrbracket \stackrel{\text{def}}{=} \bigcup \{S \mid S \subseteq \llbracket \varphi, \rho[X \mapsto S] \rrbracket\}$			
$\llbracket [p, c]\varphi, \rho \rrbracket \stackrel{\text{def}}{=} \{s \mid \forall \alpha, r, \sigma \cdot (s \xrightarrow{\alpha} r \text{ and } \text{match}(p, \alpha) = \sigma \text{ and } c\sigma \Downarrow \text{true}) \text{ implies } r \in \llbracket \varphi\sigma, \rho \rrbracket\}$			

FIGURE 2. The syntax and semantics for sHML, the safety fragment of the branching-time logic μHML [Lar90].

- *inputs*, $a?v \in \text{IACT}$ e.g., denoting an input by the system from the environment on port a carrying payload v ; and
- *outputs*, $a!v \in \text{OACT}$ originating from the system to the interacting environment on port a carrying payload v

where $\text{ACT} = \text{IACT} \cup \text{OACT}$. Systems are described as *labelled transition systems* (LTSs); these are triples $\langle \text{SYS}, \text{ACT} \cup \{\tau\}, \rightarrow \rangle$ consisting of a set of *system states*, $s, r, q \in \text{SYS}$, a set of *visible actions*, $\alpha, \beta \in \text{ACT}$, along with a distinguished silent action $\tau \notin \text{ACT}$ (where $\mu \in \text{ACT} \cup \{\tau\}$), and a *transition* relation, $\rightarrow \subseteq (\text{SYS} \times (\text{ACT} \cup \{\tau\}) \times \text{SYS})$. We write $s \xrightarrow{\mu} r$ in lieu of $(s, \mu, r) \in \rightarrow$, and $s \xrightarrow{\alpha} r$ to denote weak transitions representing $s \xrightarrow{(\tau)^*} r$ where r is called the α -derivative of s . For convenience, we use the syntax of the regular fragment of value-passing CCS [HL96] to concisely describe LTSs. Traces $t, u \in \text{ACT}^*$ range over (finite) sequences of *visible* actions. We write $s \xrightarrow{t} r$ to denote a sequence of *weak* transitions $s \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} r$ where $t = \alpha_1 \dots \alpha_n$ for some $n \geq 0$; when $t = \varepsilon$, $s \xrightarrow{\varepsilon} r$ means $s \xrightarrow{\tau^*} r$. Additionally, we represent system runs as *explicit traces* that include τ -actions, $t_\tau, u_\tau \in (\text{ACT} \cup \{\tau\})^*$ and write $s \xrightarrow{\mu_1 \dots \mu_n} r$ to denote a sequence of *strong* transitions $s \xrightarrow{\mu_1} \dots \xrightarrow{\mu_n} r$. The function $\text{sys}(t_\tau)$ returns a system that produces exclusively the sequence of actions defined by t_τ , modulo the data carried by input actions in t_τ that cannot be controlled by the receiving process. For instance, $\text{sys}(a?3.\tau.a!5)$ produces the process $a?x.\tau.a!5.\text{nil}$. We consider states in our system LTS modulo the classic notion of *strong bisimilarity* [HL96, San11] and write $s \sim r$ when states s and r are bisimilar.

The Logic. The behavioral properties we consider are described using sHML [AI99, FAI17], a subset of the value passing μHML [RH97, HL95] that uses *symbolic actions* of the form (p, c) consisting of an action pattern p and a condition c . Symbolic actions facilitate reasoning about LTSs with infinitely many actions (e.g., inputs or outputs carrying data from infinite domains). They abstract over concrete actions using *data variables* $x, y, z \in \text{DVAR}$ that

occur free in the constraint c or as binders in the pattern p . Patterns are subdivided into input $(x)?(y)$ and output $(x)!(y)$ patterns where (x) binds the information about the port on which the interaction has occurred, whereas (y) binds the payload; $\mathbf{bv}(p)$ denotes the set of binding variables in p whereas $\mathbf{fv}(c)$ represents the set of free variables in condition c . We assume a (partial) *matching function* $\text{match}(p, \alpha)$ that (when successful) returns the (smallest) substitution $\sigma : \text{DVAR} \rightarrow (\text{PORT} \cup \text{VAL})$, mapping bound variables in p to the corresponding values in α ; by replacing every occurrence (x) in p with $\sigma(x)$ we get the matched action α . The *filtering condition*, c , is evaluated wrt. the substitution returned by successful matches, written as $c\sigma \Downarrow v$ where $v \in \{\text{true}, \text{false}\}$.

Whenever a symbolic action (p, c) is *closed*, i.e., $\mathbf{fv}(c) \subseteq \mathbf{bv}(p)$, it denotes the *set* of actions $\llbracket (p, c) \rrbracket \stackrel{\text{def}}{=} \{ \alpha \mid \exists \sigma \cdot \text{match}(p, \alpha) = \sigma \text{ and } c\sigma \Downarrow \text{true} \}$. For example, we can have $\llbracket ((x)!(y), (x = \mathbf{a} \vee x = \mathbf{b}) \wedge y \geq 3) \rrbracket = \{ \mathbf{a}!3, \mathbf{b}!3, \mathbf{a}!4, \mathbf{b}!4, \mathbf{a}!5, \mathbf{b}!5, \mathbf{a}!6, \mathbf{b}!6, \dots \}$. Following standard (concrete) value-passing LTS semantics [MPW92, HL96], our systems have *no control* over the data values supplied via inputs. Accordingly, we assume a well-formedness constraint where the condition c of an input symbolic action, $((x)?(y), c)$, *cannot* restrict the values of binder y , i.e., $y \notin \mathbf{fv}(c)$. As a shorthand, whenever a condition in a symbolic action equates a bound variable to a specific value we embed the equated value within the pattern, e.g., $((x)!(y), x = \mathbf{a} \wedge y = 3)$, $((x)?(y), x = \mathbf{a})$ and $((x)?(y), x = z)$ become $(\mathbf{a}!3, \text{true})$, $(\mathbf{a}?(\mathbf{y}), \text{true})$ and $(z?(\mathbf{y}), \text{true})$ resp.; we also elide **true** conditions, and occasionally just write $(\mathbf{a}!3)$ and $(\mathbf{a}?(\mathbf{y}))$ in lieu of $(\mathbf{a}!3, \text{true})$ and $(\mathbf{a}?(\mathbf{y}), \text{true})$ when the meaning of this shorthand can be inferred from the context.

Figure 2 presents the SHML syntax for some countable set of logical variables $X, Y \in \text{LVAR}$. The construct $\bigwedge_{i \in I} \varphi_i$ describes a *compound* conjunction, $\varphi_1 \wedge \dots \wedge \varphi_n$, where $I = \{1, \dots, n\}$ is a finite set of indices. The syntax also permits recursive properties using greatest fixpoints, $\max X.\varphi$, which bind free occurrences of X in φ . The central construct is the (symbolic) universal modal operator, $[p, c]\varphi$, where the binders $\mathbf{bv}(p)$ bind the free data variables in c and φ . We occasionally use the notation $(-)$ to denote “don’t care” binders in the pattern p , whose bound values are not referenced in c and φ . We also assume that all fixpoint variables, X , are guarded by modal operators.

Formulas in SHML are interpreted over the system powerset domain where $S \in \mathcal{P}(\text{SYS})$. The semantic definition of Figure 2, $\llbracket \varphi, \rho \rrbracket$, is given for *both* open and closed formulas. It employs a valuation from logical variables to sets of states, $\rho \in (\text{LVAR} \rightarrow \mathcal{P}(\text{SYS}))$, which permits an inductive definition on the structure of the formulas; $\rho' = \rho[X \mapsto S]$ denotes a valuation where $\rho'(X) = S$ and $\rho'(Y) = \rho(Y)$ for all other $Y \neq X$. The only non-standard case is that for the universal modality formula, $[p, c]\varphi$, which is satisfied by any system that either *cannot* perform an action α that matches p while satisfying condition c , or for any such matching action α with substitution σ , its derivative state satisfies the continuation $\varphi\sigma$. We consider formulas modulo associativity and commutativity of \wedge , and unless stated explicitly, we assume *closed* formulas, i.e., without free logical and data variables. Since the interpretation of a closed φ is independent of the valuation ρ we write $\llbracket \varphi \rrbracket$ in lieu of $\llbracket \varphi, \rho \rrbracket$. A system s *satisfies* formula φ whenever $s \in \llbracket \varphi \rrbracket$, and a formula φ is *satisfiable*, when $\llbracket \varphi \rrbracket \neq \emptyset$.

We find it convenient to define the function *after*, describing how an SHML formula *evolves* in reaction to an action μ . Note that, for the case $\varphi = [p, c]\psi$, the formula returns $\psi\sigma$ when μ matches successfully the symbolic action (p, c) with σ , and **tt** otherwise, to signify a trivial satisfaction.

Definition 2.1. We define the function $after : (\text{SHML} \times \text{ACT} \cup \{\tau\}) \rightarrow \text{SHML}$ as:

$$after(\varphi, \alpha) \stackrel{\text{def}}{=} \begin{cases} \varphi & \text{if } \varphi \in \{\text{tt}, \text{ff}\} \\ after(\varphi'\{\varphi/X\}, \alpha) & \text{if } \varphi = \max X.\varphi' \\ \bigwedge_{i \in I} after(\varphi_i, \alpha) & \text{if } \varphi = \bigwedge_{i \in I} \varphi_i \\ \psi\sigma & \text{if } \varphi = [p, c]\psi \text{ and } \exists \sigma. (\text{match}(p, \alpha) = \sigma \wedge c\sigma \Downarrow \text{true}) \\ \text{tt} & \text{if } \varphi = [p, c]\psi \text{ and } \nexists \sigma. (\text{match}(p, \alpha) = \sigma \wedge c\sigma \Downarrow \text{true}) \end{cases}$$

$$after(\varphi, \tau) \stackrel{\text{def}}{=} \varphi \quad \blacksquare$$

We abuse notation and lift the $after$ function to (explicit) traces in the obvious way, i.e., $after(\varphi, t_\tau)$ is equal to $after(after(\varphi, \mu), u_\tau)$ when $t_\tau = \mu u_\tau$ and to φ when $t_\tau = \varepsilon$. Our definition of $after$ is justified vis-a-vis the semantics of Figure 2 via Proposition 2.3; it will play a role later on when defining our notion of enforcement in Section 4.

Remark 2.2. The function $after$ is well-defined due to our assumption that formulas are guarded, guaranteeing that $\varphi'\{\varphi/X\}$ has fewer top level occurrences of greatest fixpoint operators than $\max X.\varphi'$. \blacksquare

Proposition 2.3. For every system state s , formula φ and action α , if $s \in \llbracket \varphi \rrbracket$ and $s \xrightarrow{\alpha} s'$ then $s' \in \llbracket after(\varphi, \alpha) \rrbracket$. \square

Example 2.4. The *safety* property φ_1 repeatedly requires that every input request that is made on a port that is *not* \mathbf{b} , cannot be followed by another input on the same port in succession. However, following this input it allows a *single* output answer on the same port in response, followed by the logging of the serviced request by outputting a notification on a dedicated port \mathbf{b} . We note how the channel name bound to x is used to constrain sub-modalities. Similarly, values bound to y_1 and y_2 are later referenced in condition $y_3 = (\log, y_1, y_2)$.

$$\varphi_1 \stackrel{\text{def}}{=} \max X.([(x)?(y_1), x \neq \mathbf{b})]([(x)?(-)])\text{ff} \wedge [(x)!(y_2)]\varphi'_1$$

$$\varphi'_1 \stackrel{\text{def}}{=} ([(x)!(y_2)]\text{ff} \wedge [(\mathbf{b})!(y_3), y_3 = (\log, y_1, y_2)]X)$$

Consider the systems $s_{\mathbf{a}}$, $s_{\mathbf{b}}$ and $s_{\mathbf{c}}$:

$$s_{\mathbf{a}} \stackrel{\text{def}}{=} \text{rec } X.((\mathbf{a}?x.y := \text{ans}(x).\mathbf{a}!y.\mathbf{b}!(\log, x, y).X) + s_{\text{cls}})$$

$$\text{(where } s_{\text{cls}} \stackrel{\text{def}}{=} (\mathbf{b}?z.\text{if } z = \text{cls} \text{ then nil else } X))$$

$$s_{\mathbf{b}} \stackrel{\text{def}}{=} \text{rec } X.((\mathbf{a}?x.y := \text{ans}(x).\mathbf{a}!y.(\underline{\mathbf{a}}!y.\mathbf{b}!(\log, x, y).s_{\mathbf{a}} + \mathbf{b}!(\log, x, y).X) + s_{\text{cls}})$$

$$s_{\mathbf{c}} \stackrel{\text{def}}{=} \mathbf{a}?y.s_{\mathbf{a}}$$

The system $s_{\mathbf{a}}$ implements a *request-response* server that repeatedly inputs values (for some domain VAL) on port \mathbf{a} , $\mathbf{a}?x$, for which it internally computes an answer and assigns it to the data variable y , $y := \text{ans}(x)$. Subsequently, it outputs the answer on port \mathbf{a} in response to each request, $\mathbf{a}!y$, and then logs the serviced request pair of values by outputting the triple (\log, x, y) on port \mathbf{b} , $\mathbf{b}!(\log, x, y)$. It terminates whenever it receives a close request cls from port \mathbf{b} , i.e., $\mathbf{b}?z$ when $z = \text{cls}$. Systems $s_{\mathbf{b}}$ and $s_{\mathbf{c}}$ are similar to $s_{\mathbf{a}}$ but define additional behaviour: $s_{\mathbf{c}}$ requires a startup input, $\mathbf{a}?y$, before behaving as $s_{\mathbf{a}}$, whereas $s_{\mathbf{b}}$ occasionally provides a redundant (underlined) answer prior to logging a serviced request.

Using the semantics of Figure 2, one can verify that the first system satisfies our correctness property φ_1 , i.e., $s_{\mathbf{a}} \in \llbracket \varphi_1 \rrbracket$. However the second system $s_{\mathbf{b}}$ does not satisfy

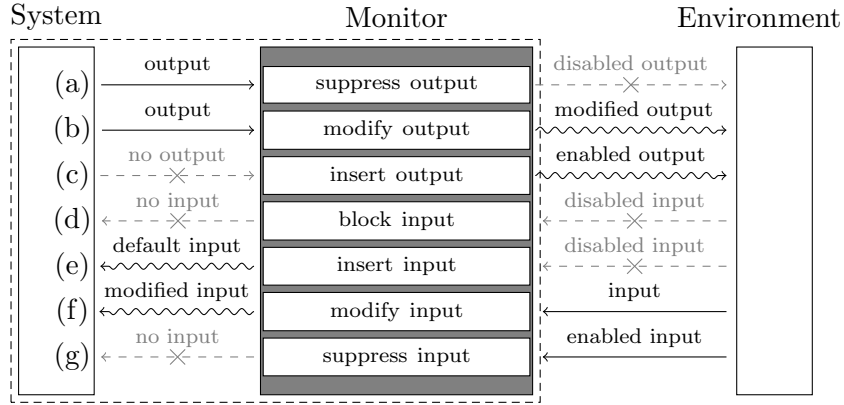


FIGURE 3. Disabling, enabling and adapting bidirectional runtime enforcement via suppressions, insertions and replacements.

this property because it can inadvertently answer twice a request, *i.e.*, $s_b \notin \llbracket \varphi_1 \rrbracket$ since we have $s_b \xrightarrow{a?v_1.a!ans(v_1).a!ans(v_1)}$ (for some value v_1). Analogously, the third system s_c violates property φ_1 because it can accept two consecutive inputs on port a (without answering the preliminary request first), *i.e.*, $s_c \notin \llbracket \varphi_1 \rrbracket$ since $s_c \xrightarrow{a?v_1.a?v_2}$ (for any pair of values v_1 and v_2). ■

3. A BIDIRECTIONAL ENFORCEMENT MODEL

Bidirectional enforcement seeks to transform the entire (visible) behaviour of the SuS in terms of output actions (instigated by the SuS itself, which in turn controls the payload values being communicated) and input actions (originating from the interacting environment which chooses the payload values); this contrasts with unidirectional approaches that only modify output traces. In this richer setting, it helps to differentiate between the transformations performed by the monitor (*i.e.*, insertions, suppressions and replacements), and the way they can be used to affect the resulting behaviour of the composite system. In particular, we say that:

- an action that can be performed by the SuS has been *disabled* when it is no longer visible in the resulting composite system (consisting of the SuS and the monitor);
- an action is *enabled* when the composite system can execute it while the SuS cannot;
- an action is *adapted* when either its payload differs from that of the composite system, or when the action is rerouted through a different port.

We argue that implementing action enabling, disabling and adaptation differs according to whether the action is an input or an output; see Figure 3. Enforcing actions instigated by the SuS—such as outputs—is more straightforward. Figure 3(a), (b) and (c) *resp.* state that disabling an output can be achieved by suppressing it, adapting an output amounts to replacing the payload or redirecting it to a different port, whereas output enabling can be attained via an insertion. However, enforcing actions instigated by the environment such as inputs is harder. In Figure 3(d), we propose to *disable* an input by concealing the input port. Since this may block the SuS from progressing, the instrumented monitor may additionally

Syntax

$$m, n \in \text{TRN} ::= (p, c, p').m \mid \sum_{i \in I} m_i \text{ (} I \text{ is a finite index set)} \mid \text{rec } X.m \mid X$$

Dynamics

$$\begin{array}{c} \text{ESEL} \frac{m_j \xrightarrow{\gamma \blacktriangleright \gamma'} n_j}{\sum_{i \in I} m_i \xrightarrow{\gamma \blacktriangleright \gamma'} n_j} \quad j \in I \qquad \text{EREC} \frac{m \{ \text{rec } X.m / X \} \xrightarrow{\gamma \blacktriangleright \gamma'} n}{\text{rec } X.m \xrightarrow{\gamma \blacktriangleright \gamma'} n} \\ \\ \text{ETRN} \frac{\text{match}(p, \gamma) = \sigma \quad c\sigma \Downarrow \text{true} \quad \gamma' = \pi\sigma}{(p, c, \pi).m \xrightarrow{\gamma \blacktriangleright \gamma'} m\sigma} \end{array}$$

Instrumentation

$$\begin{array}{c} \text{BITRNO} \frac{s \xrightarrow{\text{b!}w} s' \quad m \xrightarrow{(\text{b!}w) \blacktriangleright (\text{a!}v)} n}{m[s] \xrightarrow{\text{a!}v} n[s']} \qquad \text{BITRNI} \frac{m \xrightarrow{(\text{a?}v) \blacktriangleright (\text{b?}w)} n \quad s \xrightarrow{\text{b?}w} s'}{m[s] \xrightarrow{\text{a?}v} n[s']} \\ \\ \text{BIDISO} \frac{s \xrightarrow{\text{a!}v} s' \quad m \xrightarrow{(\text{a!}v) \blacktriangleright \bullet} n}{m[s] \xrightarrow{\tau} n[s']} \qquad \text{BIDISI} \frac{m \xrightarrow{\bullet \blacktriangleright (\text{a?}v)} n \quad s \xrightarrow{\text{a?}v} s'}{m[s] \xrightarrow{\tau} n[s']} \\ \\ \text{BIENO} \frac{m \xrightarrow{\bullet \blacktriangleright (\text{a!}v)} n}{m[s] \xrightarrow{\text{a!}v} n[s]} \qquad \text{BIENI} \frac{m \xrightarrow{(\text{a?}v) \blacktriangleright \bullet} n}{m[s] \xrightarrow{\text{a?}v} n[s]} \qquad \text{BIASY} \frac{s \xrightarrow{\tau} s'}{m[s] \xrightarrow{\tau} m[s']} \\ \\ \text{BIDEF} \frac{s \xrightarrow{\text{a!}v} s' \quad m \xrightarrow{\text{a!}v} \quad \forall \text{b} \in \text{PORT}, w \in \text{VAL} \cdot m \xrightarrow{\bullet \blacktriangleright \text{b!}w}}{m[s] \xrightarrow{\text{a!}v} \text{id}[s']}$$

FIGURE 4. A bidirectional instrumentation model for enforcement monitors.

insert a default input to unblock the system waiting to input on the channel used for the insertion, Figure 3(e), in cases where the environment fails to provide the corresponding output. Input *adaptation*, Figure 3(f), is also attained via a *replacement*, albeit applied in the opposite direction to the output case. Inputs can also be *enabled* whenever the SuS is unable to carry them out, Figure 3(g), by having the monitor accept the input in question and then *suppress* it. Note that, from the perspective of the environment, the input would still be effected.

Figure 4 presents an operational model for the bidirectional instrumentation proposal of Figure 3 in terms of (symbolic) transducers. A variant of these transducers was originally introduced in [ACFI18] for unidirectional enforcement. Transducers, $m, n \in \text{TRN}$, are monitors that define *symbolic transformation triples*, (p, c, π) , consisting of an action *pattern* p , *condition* c , and a *transformation action* π . Conceptually, the action pattern and condition determine the range of system (input or output) actions upon which the transformation should be applied, while the transformation action specifies the transformation that should be applied. The symbolic transformation pattern p is an extended version of those definable

in symbolic actions, that may also include \bullet ; when $p = \bullet$, it means that the monitor can act independently from the system to insert the action specified by the transformation action. Transformation actions are possibly open actions (*i.e.*, actions with possibly free variable such as $x?v$ or $a!x$) or the special action \bullet ; the latter represents the suppression of the action specified by p . We assume a well-formedness constraint where, for every $(p, c, \pi).m$, p and π cannot both be \bullet , and when neither is, they are of the *same* type *i.e.*, an input (*resp.* output) pattern and action. Examples of well-formed symbolic transformations are:

- $(\bullet, \text{true}, a?v)$, inserting an input on port a with value v ;
- $((x)!(y), y \geq 5, \bullet)$, suppressing an output action carrying a payload that is greater or equal to 5; and
- $((x)!(y), x = b, a!y)$, redirecting (*i.e.*, adapting) outputs on port b carrying the payload y (learnt dynamically at runtime) to port a .

The monitor transition rules in Figure 4 assume closed terms, *i.e.*, every *transformation-prefix transducer* of the form $(p, c, \pi).m$ must obey the closure constraint stating that $(\mathbf{fv}(c) \cup \mathbf{fv}(\pi) \cup \mathbf{fv}(m)) \subseteq \mathbf{bv}(p)$. A similar closure requirement is assumed for recursion variables X and $\text{rec } X.m$. Each transformation-prefix transducer yields an LTS with labels of the form $\gamma \blacktriangleright \gamma'$, where $\gamma, \gamma' \in (\text{ACT} \cup \{\bullet\})$. Intuitively, transition $m \xrightarrow{\gamma \blacktriangleright \gamma'} n$ denotes the way that a transducer in state m *transforms* the action γ into γ' while transitioning to state n . The transducer action $\alpha \blacktriangleright \beta$ represents the *replacement* of α by β , $\alpha \blacktriangleright \alpha$ denotes the *identity* transformation, whereas $\alpha \blacktriangleright \bullet$ and $\bullet \blacktriangleright \alpha$ respectively denote the *suppression* and *insertion* transformations of action α . The key transition rule in Figure 4 is ETRN. It states that the transformation-prefix transducer $(p, c, \pi).m$ transforms action γ into a (potentially) different action γ' and reduces to state $m\sigma$, whenever γ matches pattern p , *i.e.*, $\text{match}(p, \gamma) = \sigma$, and satisfies condition c , *i.e.*, $c\sigma \Downarrow \text{true}$. Action γ' results from instantiating the free variables in π as specified by σ , *i.e.*, $\gamma' = \pi\sigma$. The remaining rules for selection (ESEL) and recursion (EREC) are standard. We employ the shorthand notation $m \xrightarrow{\gamma} n$ to mean $\nexists \gamma', m'$ such that $m \xrightarrow{\gamma \blacktriangleright \gamma'} m'$. Moreover, for the semantics of Figure 4, we can encode the identity transducer/monitor, id , as follows

$$\text{id} \stackrel{\text{def}}{=} \text{rec } Y.((x)!(y), \text{true}, x!y).Y + ((x)?(y), \text{true}, x?v).Y. \quad (3.1)$$

When instrumented with any arbitrary system, the identity monitor id leaves its behaviour unchanged. As a shorthand notation, we write $(p, c).m$ instead of $(p, c, \pi).m$ when all the binding occurrences (x) in p correspond to free occurrences x in π , thus denoting an identity transformation. Similarly, we elide c whenever $c = \text{true}$.

The first contribution of this work lies in the new *instrumentation relation* of Figure 4, linking the behaviour of the SuS s with that of a monitor m : the term $m[s]$ denotes their composition as a *monitored system*. Crucially, the instrumentation rules in Figure 4 give us a semantics in terms of an LTS over the actions $\text{ACT} \cup \{\tau\}$, in line with the LTS semantics of the SuS. Following Figure 3(b), rule BTRNO states that if the SuS transitions with an output $b!w$ to s' and the transducer can *replace* it with $a!v$ and transition to n , the *adapted* output can be externalised so that the composite system $m[s]$ transitions over $a!v$ to $n[s']$. Rule BIDISO states that if s performs an output $a!v$ that the monitor *can suppress*, the instrumentation withholds this output and the composite system silently transitions; this amounts to action *disabling* as outlined in Figure 3(a). Rule BIENO is dual, and it *enables* the output $a!v$ on the SuS as outlined in Figure 3(c): it augments the composite system $m[s]$ with an output $a!v$ whenever m can *insert* $a!v$, independently of the behaviour of s . Rules

BIDISO, BITRNO and BIENO therefore correspond to items (a), (b) and (c) in Figure 3 respectively.

Rule BIDEF is analogous to standard rules for premature monitor termination [Fra21, FAI17, Fra17, AAFI18a], and accounts for underspecification of transformations. We, however, restrict defaulting (termination) to output actions performed by the SuS exclusively, *i.e.*, a monitor only defaults to `id` when it cannot react to or enable a system output. By forbidding the monitor from defaulting upon unspecified inputs, the monitor is able to *block* them from becoming part of the composite system's behaviour. Hence, any input that the monitor is unable to react to, *i.e.*, $m \xrightarrow{a?v\gamma}$, is considered as being *invalid and blocked* by default. This technique is thus used to implement Figure 3(d). To avoid disabling valid inputs unnecessarily, the monitor must therefore explicitly define symbolic transformations that cover *all* the valid inputs of the SuS. Note, that rule BIASY still allows the SuS to silently transition independently of m . Following Figure 3(f), rule BITRNI adapts inputs, provided the SuS can accept the adapted input. Similarly, rule BIENI *enables* an input on a port a as described in Figure 3(g): the composite system accepts the input while suppressing it from the SuS. Rule BIDISI allows the monitor to generate a default input value v and forward it to the SuS on a port a , thereby unblocking it whenever the environment is unable to provide the corresponding output on channel a (carrying v); from the environment's perspective, the composite system silently transitions to some state, following Figure 3(e). It is worth comparing rule BIDISI with the other instrumentation rule BIENO discussed earlier, since they both handle outputs inserted by the monitor. In the case of rule BIENO, whenever the monitor inserts an output to be consumed *by the environment*, this is expressed at the level of the composite system as an external output (see conclusion of rule BIENO) since, in our LTS, the actions represent the interaction between the (composite) system and the environment. Contrastingly, whenever the monitor inserts an output to be input *by the SuS*, then this is expressed at the level of the composite system as a silent action (see conclusion of rule BIDISI) since no interaction occurs between the (composite) system and the environment. We conclude our discussion of the instrumentation rules in Figure 4 by remarking that rules BIDISI, BITRNI and BIENI respectively implement items (e), (f) and (g) of Figure 3.

Definition 3.1. We call disabling monitors/transducers those monitors that only perform disabling actions. The same applies to enabling and adapting monitors/transducers. ■

Example 3.2. Consider the following action disabling transducer m_d , that repeatedly disables every output performed by the system via the branch $((-)!(-), \bullet).Y$. In addition, it limits inputs to those on port b via the input branch $(b?(-)).Y$; inputs on other ports are disabled since none of the relevant instrumentation rules in Figure 4 can be applied.

$$m_d \stackrel{\text{def}}{=} \text{rec } Y.(b?(-)).Y + ((-)!(-), \bullet).Y$$

Recall the two systems below from Example 2.4:

$$s_b \stackrel{\text{def}}{=} \text{rec } X.((a?x.y := \text{ans}(x).a!y.(a!y.b!(\log, x, y).s_a + b!(\log, x, y).X)) + s_{\text{cls}})$$

$$s_c \stackrel{\text{def}}{=} a?y.s_a$$

where

$$s_a \stackrel{\text{def}}{=} \text{rec } X.((a?x.y := \text{ans}(x).a!y.b!(\log, x, y).X) + s_{\text{cls}}) \quad \text{and}$$

$$s_{\text{cls}} \stackrel{\text{def}}{=} (b?z.\text{if } z = \text{cls} \text{ then nil else } X)$$

When instrumented with the system s_c , monitor m_d blocks its initial input, *i.e.*, we have $m_d[s_c] \not\rightarrow^\alpha$ for any α . In the case of s_b , the composite system $m_d[s_b]$ can only input requests on port b , such as the termination request $m_d[s_b] \xrightarrow{b?cls} m_d[nil]$.

$$\begin{aligned} m_{dt} &\stackrel{\text{def}}{=} \text{rec } X.(((x)?(y_1), x \neq b).(((x_1)?(-), x_1 \neq x).id + (x!(y_2)).m'_{dt}) + (b?(-)).id) \\ m'_{dt} &\stackrel{\text{def}}{=} (x!(-), \bullet).m_d + ((-)?(-)).id + (b!(y_3), y_3 = (\log, y_1, y_2)).X \end{aligned}$$

By defining branch $(b?(-)).id$, the more elaborate monitor m_{dt} (above) allows the SuS to immediately input on port b (possibly carrying a termination request). At the same time, the branch prefixed by $((x)?(y_1), x \neq b)$ permits the SuS to input the first request via any port $x \neq b$, subsequently blocking inputs on the same port x (without deterring inputs on other ports) via the input branch $((x_1)?(-), x_1 \neq x).id$. In conjunction to this branch, m_{dt} defines another branch $(x!(y_2)).m'_{dt}$ to allow outputs on the port bound to variable x . The continuation monitor m'_{dt} then defines the suppression branch $(x!(-), \bullet).m_d$ by which it disables any *redundant* response that is output following the first one. Since it also defines branches $(b!(y_3), y_3 = (\log, y_1, y_2)).X$ and $((-)?(-)).id$, it does not affect log events or further inputs that occur immediately after the first response.

When instrumented with system s_c from Example 2.4, m_{dt} allows the composite system to perform the first input but then blocks the second one, permitting only input requests on channel b , *e.g.*,

$$m_{dt}[s_c] \xrightarrow{a?v} \cdot \xrightarrow{b?cls} id[nil].$$

It also disables the first redundant response of system s_b while transitioning to m_d , which proceeds to suppress every subsequent output (including log actions) while blocking every other port except b , *i.e.*,

$$m_{dt}[s_b] \xrightarrow{a?v} \cdot \xrightarrow{a!w} \cdot \xrightarrow{\tau} m_d[b!(\log, v, w).s_a] \xrightarrow{\tau} m_d[s_a] \not\rightarrow^{a?v}$$

(for every port a where $a \neq b$ and any value v). Rule IDEF allows it to default when handling unspecified outputs, *e.g.*, for system $b!(\log, v, w).s_a$ the composite system can still perform the logging output, *i.e.*,

$$m_{dt}[b!(\log, v, w).s_a] \xrightarrow{b!(\log, v, w)} id[s_a].$$

Consider one further monitor, defined below:

$$\begin{aligned} m_{det} &\stackrel{\text{def}}{=} \text{rec } X.(((x)?(y_1), x \neq b).m'_{det} + (b?(-)).id) \\ m'_{det} &\stackrel{\text{def}}{=} \text{rec } Y_1.(\bullet, x?v_{\text{def}}).\underline{Y_1} + (x!(y_2)).m''_{det} + ((x_1)?(-), x_1 \neq x).id \\ m''_{det} &\stackrel{\text{def}}{=} \text{rec } Y_2.(\underline{(x!(-), x \neq b, \bullet)}.Y_2 + (b!(y_3), y_3 = (\log, y_1, y_2)).X + ((-)?(-)).id) \end{aligned}$$

Monitor m_{det} (above) behaves similarly to m_{dt} but instead employs a loop of suppressions (underlined in m''_{det}) to disable further responses until a log or termination input is made. When composed with s_b , it permits the log action to go through:

$$m_{det}[s_b] \xrightarrow{a?v} \cdot \xrightarrow{a!w} \cdot \xrightarrow{\tau} m''_{det}[b!(\log, v, w).s_b] \xrightarrow{b!(\log, v, w)} m_{det}[s_b].$$

m_{det} also defines a branch prefixed by the insertion transformation $(\bullet, x?v_{\text{def}})$ (underlined in m'_{det}) where v_{def} is a default input domain value. This permits the instrumentation to silently unblock the SuS when this is waiting for a request following an unanswered one. In fact, when instrumented with s_c , m_{det} not only forbids invalid input requests,

but it also (internally) unblocks s_c by supplying the required input via the added insertion branch. This allows the composite system to proceed, as shown below (where $s'_a \stackrel{\text{def}}{=} y := \text{ans}(v_{\text{def}}).a!y.b!(\log, v_{\text{def}}, y).s_a$):

$$\begin{aligned} m_{\text{det}}[s_c] &\xrightarrow{a?v} \text{rec } Y.((\bullet, a?v_{\text{def}}).Y + (a!(y_2)).m''_{\text{det}} + (b?(-)).\text{id})[s_a] \\ &\xrightarrow{\tau} \text{rec } Y.((\bullet, a?v_{\text{def}}).Y + (a!(y_2)).m''_{\text{det}} + (b?(-)).\text{id})[s'_a] \\ &\xrightarrow{a!\text{ans}(v_{\text{def}}).b!(\log, v_{\text{def}}, y)} m_{\text{det}}[s_a] \quad \blacksquare \end{aligned}$$

Although in this paper we mainly focus on action disabling monitors, using our model one can also define action enabling and adaptation monitors.

Example 3.3. Consider now the transducers m_e and m_a below:

$$\begin{aligned} m_e &\stackrel{\text{def}}{=} ((x)?(y), x \neq b, \bullet).(\bullet, x!\text{ans}(y)).(\bullet, b!(\log, y, \text{ans}(y))).\text{id} \\ m_a &\stackrel{\text{def}}{=} \text{rec } X.(b?(y), a?y).X + (a!(y), b!y).X. \end{aligned}$$

Once instrumented, m_e first uses a suppression to enable an input on any port $x \neq b$ (but then gets discarded). It then automates a response by inserting an answer followed by a log action. Concretely, when composed with the systems $r \in \{s_b, s_c\}$ from Example 2.4 (restated in Example 3.2), the execution of the composite system can only start as follows, for some channel name $c \neq b$, values v and $w = \text{ans}(v)$:

$$m_e[r] \xrightarrow{c?v} (\bullet, c!w).(\bullet, b!(\log, v, w)).\text{id}[r] \xrightarrow{c!w} (\bullet, b!(\log, v, w)).\text{id}[r] \xrightarrow{b!(\log, v, w)} \text{id}[r].$$

By contrast, m_a uses action adaptation to redirect the inputs and outputs from the SuS through port b : it allows the composite system to exclusively input values on port b forwarding them to the SuS on port a , and dually allowing outputs from the SuS on port a to reroute them to port b . As a result, from an external viewpoint, the resulting composite system can only be seen to communicate on port b with its environment. For instance, for the systems s_c and s_b restated earlier, we can observe the following behaviour:

$$\begin{aligned} m_a[s_c] &\xrightarrow{b?v_1} m_a[s_a] \xrightarrow{b?v_2.b!w_2.b!(\log, v_2, w_2)} m_a[s_a] \quad \text{and} \\ m_a[s_b] &\xrightarrow{b?v_1.b!w_1.b!(\log, v_1, w_1)} m_a[s_b]. \quad \blacksquare \end{aligned}$$

4. ENFORCEMENT

We are concerned with extending the enforceability result obtained in prior work [ACFI18] to the extended setting of bidirectional enforcement. The *enforceability* of a logic rests on the relationship between the semantic behaviour specified by the logic on the one hand, and the ability of the operational mechanism (that of Section 3 in this case) to enforce the specified behaviour on the other. This is captured by the predicate “(monitor) m adequately enforces (property) φ ” in Definition 4.1 below. In fact, the definitions of formula and logic enforceability in Definition 4.1 are parametric with respect to the precise meaning of such a predicate. In what follows, we will explore the design space for formalising this predicate.

Definition 4.1 (Enforceability [ACFI18]). A formula φ is *enforceable* iff there *exists* a transducer m such that m *adequately enforces* φ . A logic \mathcal{L} is *enforceable* iff *every* formula $\varphi \in \mathcal{L}$ is *enforceable*. ■

Since we have limited control over the SuS that a monitor is composed with, “ m adequately enforces φ ” should hold for *any* (instrumentable) system. In [ACFI18] we stipulate that any notion of adequate enforcement should at least entail soundness.

Definition 4.2 (Sound Enforcement [ACFI18]). Monitor m *soundly enforces* a formula φ , denoted as $\text{senf}(m, \varphi)$, iff, whenever φ is satisfiable, i.e., $\llbracket \varphi \rrbracket \neq \emptyset$, then for every state $s \in \text{Sys}$, it is the case that $m[s] \in \llbracket \varphi \rrbracket$. ■

Example 4.3. Although showing that a monitor soundly enforces a formula should consider *all* systems, we give an intuition based on $s_{\mathbf{a}}$, $s_{\mathbf{b}}$, $s_{\mathbf{c}}$ for formula φ_1 from Example 2.4 (restated below) where $s_{\mathbf{a}} \in \llbracket \varphi_1 \rrbracket$ (hence $\llbracket \varphi_1 \rrbracket \neq \emptyset$) and $s_{\mathbf{b}}, s_{\mathbf{c}} \notin \llbracket \varphi_1 \rrbracket$.

$$\begin{aligned}\varphi_1 &\stackrel{\text{def}}{=} \max X. [(x)?(y_1), x \neq \mathbf{b}] ([(x)?(-)] \text{ff} \wedge [(x!(y_2))] \varphi'_1) \\ \varphi'_1 &\stackrel{\text{def}}{=} ([(x!(-))] \text{ff} \wedge [(\mathbf{b}!(y_3), y_3 = (\log, y_1, y_2))] X)\end{aligned}$$

Recall the transducers $m_{\mathbf{e}}$, $m_{\mathbf{a}}$, $m_{\mathbf{d}}$, $m_{\mathbf{dt}}$ and $m_{\mathbf{det}}$ from Examples 3.2 and 3.3, restated below:

$$\begin{aligned}m_{\mathbf{e}} &\stackrel{\text{def}}{=} ((x)?(y), x \neq \mathbf{b}, \bullet).(\bullet, x! \text{ans}(y)).(\bullet, \mathbf{b}!(\log, y, \text{ans}(y))).\text{id} \\ m_{\mathbf{a}} &\stackrel{\text{def}}{=} \text{rec } X. (\mathbf{b}?(y), \mathbf{a}?y).X + (\mathbf{a}!(y), \mathbf{b}!y).X \\ m_{\mathbf{d}} &\stackrel{\text{def}}{=} \text{rec } Y. (\mathbf{b}?(-)).Y + ((-)!(-), \bullet).Y \\ m_{\mathbf{dt}} &\stackrel{\text{def}}{=} \text{rec } X. (((x)?(y_1), x \neq \mathbf{b}).(((x_1)?(-), x_1 \neq x).\text{id} + (x!(y_2)).m'_{\mathbf{dt}}) + (\mathbf{b}?(-)).\text{id}) \\ m_{\mathbf{det}} &\stackrel{\text{def}}{=} \text{rec } X. (((x)?(y_1), x \neq \mathbf{b}).m'_{\mathbf{det}} + (\mathbf{b}?(-)).\text{id})\end{aligned}$$

where

$$\begin{aligned}m'_{\mathbf{dt}} &\stackrel{\text{def}}{=} (x!(-), \bullet).m_{\mathbf{d}} + ((-)?(-)).\text{id} + (\mathbf{b}!(y_3), y_3 = (\log, y_1, y_2)).X \\ m''_{\mathbf{det}} &\stackrel{\text{def}}{=} \text{rec } Y_1. (\bullet, x?v_{\text{def}}).Y_1 + (x!(y_2)).m''_{\mathbf{det}} + ((x_1)?(-), x_1 \neq x).\text{id} \\ m'''_{\mathbf{det}} &\stackrel{\text{def}}{=} \text{rec } Y_2. ((x!(-), x \neq \mathbf{b}, \bullet).Y_2 + (\mathbf{b}!(y_3), y_3 = (\log, y_1, y_2)).X + ((-)?(-)).\text{id})\end{aligned}$$

When assessing their soundness in relation to the property φ_1 , we have that:

- $m_{\mathbf{e}}$ is *unsound* for φ_1 . When composed with $s_{\mathbf{b}}$, the resulting monitored system produces two consecutive output replies (namely those underlined in the trace $t_{\mathbf{e}}^1$ below), thus meaning that the composite system violates the property in question, i.e., $m_{\mathbf{e}}[s_{\mathbf{b}}] \notin \llbracket \varphi_1 \rrbracket$. More concretely, we have

$$m_{\mathbf{e}}[s_{\mathbf{b}}] \xrightarrow{t_{\mathbf{e}}^1} \text{id}[s_{\mathbf{b}}] \quad \text{where } t_{\mathbf{e}}^1 \stackrel{\text{def}}{=} c?v_1.c! \text{ans}(v_1).\mathbf{b}!(\log, v_1, \text{ans}(v_1)).\mathbf{a}?v_2.\mathbf{a}!w_2.\mathbf{a}!w_2.$$

Similarly, the system $s_{\mathbf{c}}$ instrumented with the transducer $m_{\mathbf{e}}$ also violates property φ_1 , i.e., $m_{\mathbf{e}}[s_{\mathbf{c}}] \notin \llbracket \varphi_1 \rrbracket$, since the $m_{\mathbf{e}}[s_{\mathbf{c}}]$ executes the erroneous trace with two consecutive inputs on port \mathbf{a} (underlined), $c?v_1.c! \text{ans}(v_1).\mathbf{b}!(\log, v_1, \text{ans}(v_1)).\mathbf{a}?w_2.\mathbf{a}?w_3$. This demonstrates that $m_{\mathbf{e}}[s_{\mathbf{c}}]$ can still input two consecutive requests on port \mathbf{a} (underlined). Either one of these counterexamples *disproves* $\text{senf}(m_{\mathbf{e}}, \varphi_1)$.

- Monitor $m_{\mathbf{a}}$ turns out to be *sound* for φ_1 because once instrumented, the resulting composite system is adapted to only interact on port \mathbf{b} . In fact, we have $\{m_{\mathbf{a}}[s_{\mathbf{a}}], m_{\mathbf{a}}[s_{\mathbf{b}}], m_{\mathbf{a}}[s_{\mathbf{c}}]\} \subseteq \llbracket \varphi_1 \rrbracket$. The other monitors $m_{\mathbf{d}}$, $m_{\mathbf{dt}}$ and $m_{\mathbf{det}}$ are also *sound* for φ_1 . Whereas, monitor $m_{\mathbf{d}}$ prevents the violation of φ_1 by also blocking all input ports except \mathbf{b} , the transducers $m_{\mathbf{dt}}$ and $m_{\mathbf{det}}$ achieve the same goal by disabling the invalid consecutive requests and answers that occur on any port except \mathbf{b} . ■

By itself, sound enforcement is a weak criterion because it does not regulate the *extent* to which enforcement is applied. More specifically, although $m_{\mathbf{d}}$ from Example 3.2 is sound, it needlessly modifies the behaviour of $s_{\mathbf{a}}$ even though $s_{\mathbf{a}}$ satisfies φ_1 : by blocking the initial input of $s_{\mathbf{a}}$, $m_{\mathbf{d}}$ causes it to block indefinitely. The requirement that a monitor should not modify the behaviour of a system that satisfies the property being enforced can be formalised using a transparency criterion.

Definition 4.4 (Transparent Enforcement [ACFI18]). A monitor m *transparently* enforces a formula φ , $\text{tenf}(m, \varphi)$, iff for *all* $s \in \text{SYS}$, $s \in \llbracket \varphi \rrbracket$ implies $m[s] \sim s$. \square

Example 4.5. As argued earlier, $s_{\mathbf{a}}$ suffices to disprove $\text{tenf}(m_{\mathbf{d}}, \varphi_1)$. Monitor $m_{\mathbf{a}}$ from Example 3.3 also breaches Definition 4.4: although $s_{\mathbf{a}} \in \llbracket \varphi_1 \rrbracket$, we have $m_{\mathbf{a}}[s_{\mathbf{a}}] \not\sim s_{\mathbf{a}}$ since for any value v and w , $s_{\mathbf{a}} \xrightarrow{a?v} \cdot$ but for any value v we can *never* have $m_{\mathbf{a}}[s_{\mathbf{a}}] \xrightarrow{a?v} \cdot$. By contrast, monitors $m_{\mathbf{dt}}$ and $m_{\mathbf{det}}$ turn out to satisfy Definition 4.4, since they only intervene when it becomes apparent that a violation will occur. For instance, they only disable inputs on a specific port, as a precaution, following an unanswered request on the same port, and they only disable the redundant responses that are produced after the first response to a request. \blacksquare

It turns out that, by some measures, Definition 4.4 is still a relatively weak requirement since it only limits transparency requirements to *well-behaved* systems, *i.e.*, those that satisfy the property in question, and disregards enforcement behaviour for systems that violate this property. For instance, consider monitor $m_{\mathbf{dt}}$ from Example 3.2 (restated in Example 4.3) and system $s_{\mathbf{b}}$ from Example 2.4 (restated in Example 3.2). At runtime, $s_{\mathbf{b}}$ can exhibit the following invalid behaviour:

$$s_{\mathbf{b}} \xrightarrow{t_1} \mathbf{b}!(\log, v, w).s_{\mathbf{a}} \quad \text{where } t_1 \stackrel{\text{def}}{=} a?v.a!w \text{ for some appropriate pair of values } v, w.$$

In order to rectify this violating behaviour *wrt.* formula φ_1 , it suffices to use a monitor that disables *one* of the responses in t_1 , *i.e.*, $a!w$. Following this disabling, no further modifications are required since the SuS reaches a state that does not violate the remainder of the formula φ_1 , *i.e.*, $\mathbf{b}!(\log, v, w).s_{\mathbf{a}} \in \llbracket \text{after}(\varphi_1, t'_1) \rrbracket$ where $t'_1 \stackrel{\text{def}}{=} a?v.a!w$. However, when instrumented with $m_{\mathbf{dt}}$, this monitor does not only disable the invalid response, namely $m_{\mathbf{dt}}[s_{\mathbf{b}}] \xrightarrow{a?v.a!w} m_{\mathbf{d}}[\mathbf{b}!(\log, v, w).s_{\mathbf{a}}]$, but subsequently disables every other action by reaching $m_{\mathbf{d}}$, $m_{\mathbf{d}}[\mathbf{b}!(\log, v, w).s_{\mathbf{a}}] \xrightarrow{\tau} m_{\mathbf{d}}[s_{\mathbf{a}}]$. To this end, we introduce the novel requirement of *eventual transparency*.

Definition 4.6 (Eventually Transparent Enforcement). Monitor m enforces property φ in an eventually transparent way, $\text{eventf}(m, \varphi)$, iff for all systems s, s' , traces t and monitors m' , $m[s] \xrightarrow{t} m'[s']$ and $s' \in \llbracket \text{after}(\varphi, t) \rrbracket$ imply $m'[s'] \sim s'$. \blacksquare

Example 4.7. We have already argued why $m_{\mathbf{dt}}$ (restated in Example 4.3) does not adhere to eventual transparency via the counterexample $s_{\mathbf{b}}$. This is not the case for $m_{\mathbf{det}}$ (also restated in Example 4.3). Although the universal quantification over all systems and traces make it hard to prove this property, we get an intuition of why this is the case from the system $s_{\mathbf{b}}$. More concretely, when

$$m_{\mathbf{det}}[s_{\mathbf{b}}] \xrightarrow{a?v_1.a!w_1} \cdot \xrightarrow{\tau} m''_{\mathbf{det}}[\mathbf{b}!(\log, v_1, w_1).s_{\mathbf{a}}]$$

we have

$$\mathbf{b}!(\log, v_1, w_1).s_{\mathbf{a}} \in \llbracket \text{after}(\varphi_1, \mathbf{a}?v_1.\mathbf{a}!w_1) \rrbracket$$

since

$$\text{after}(\varphi_1, \mathbf{a}?v_1.\mathbf{a}!w_1) = (((x_3)!(-), x_3=\mathbf{a})\text{ff} \wedge (((x_4)!(y_3), x_4=\mathbf{b} \wedge y_3=(\log, v_1, w_1)))\varphi_1)$$

and, moreover, $m''_{\text{det}}[\mathbf{b}!(\log, v_1, w_1).s_{\mathbf{a}}] \sim \mathbf{b}!(\log, v_1, w_1).s_{\mathbf{a}}$. ■

Corollary 4.8. *For all monitors $m \in \text{TRN}$ and properties $\varphi \in \text{sHML}$, $\text{eventf}(m, \varphi)$ implies $\text{tenf}(m, \varphi)$.* □

Along with Definition 4.2 (soundness), Definition 4.6 (eventual transparency) makes up our definition for “ m (adequately) enforces φ ”. From Corollary 4.8, it follows that is definition is stricter than the one given in [ACFI18].

Definition 4.9 (Adequate Enforcement). A monitor m (adequately) *enforces* property φ , denoted as $\text{enf}(n, \varphi)$, iff it adheres to (i) *soundness*, Definition 4.2, and (ii) *eventual transparency*, Definition 4.6. ■

5. SYNTHESISING ACTION DISABLING MONITORS

Although Definition 4.1 (instantiated with Definition 4.9) enables us to rule out erroneous monitors that purport to enforce a property, the *universal quantifications* over all systems in Definitions 4.2 and 4.6 make it difficult to prove that a monitor does indeed enforce a property correctly in a bidirectional setting (disproving, however, is easier). Establishing that a formula is enforceable, Definition 4.9, involves a further existential quantification over a monitor that enforces it correctly; put differently, in order to show that a formula is *not* enforceable, amounts to another universal quantification, this time over all possible monitors. Moreover, establishing the enforceability of a logic entails yet another universal quantification, on all the formulas in the logic. In many cases (including ours), the sets of systems, monitors and formulas are infinite.

We address these problems through an *automated synthesis procedure* that produces an enforcement monitor from a safety μHML formula, expressed in a syntactic fragment of sHML. This fragment, called sHML_{nf} , has already been used to establish enforceability results in a uni-directional setting [ACFI18] and is the source logic employed by the tool detectEr^1 [AAA⁺21, AAA⁺22, AEF⁺22] used to verify the correctness of concurrent systems written in Erlang [AAFI21] and Elixir [BAF21]; it also coincides with sHML in the regular setting [AAF⁺20]. We show that the synthesised monitors are correct, according to Definition 4.9. For a unidirectional setting, it has been shown that monitors that only administer *omissions* are expressive enough to enforce *safety properties* [LBW05, FFM12, vHRF17, ACFI18]. Analogously, for our bidirectional case, we restrict ourselves to action disabling monitors and show that they can enforce *any* property expressed in terms of this sHML fragment.

Our synthesis procedure is compositional, meaning that the monitor synthesis of a composite formula is defined in terms of the enforcement monitors generated from its constituent sub-formulas. Compositionality simplifies substantially our correctness analysis of the generated monitors (e.g., we can use standard inductive proof techniques). The choice

¹<https://duncanatt.github.io/detector/>

of the logical fragment, i.e., SHML_{nf} , facilitates this compositional definition. An automated procedure to translate an SHML formula with symbolic actions where the scope of the data binders is limited to the immediate symbolic action condition, into a corresponding SHML_{nf} one (with the same semantic meaning) is given in [ACFI18, AAF⁺20].

Definition 5.1 (SHML normal form). The set of normalised SHML formulas is generated by the following grammar (where² $|I| \geq 1$):

$$\varphi, \psi \in \text{SHML}_{\text{nf}} ::= \text{tt} \quad | \quad \text{ff} \quad | \quad \bigwedge_{i \in I} [p_i, c_i] \varphi_i \quad | \quad X \quad | \quad \max X. \varphi .$$

In addition, SHML_{nf} formulas are required to satisfy the following conditions:

- (1) Every branch in $\bigwedge_{i \in I} [p_i, c_i] \varphi_i$, must be *disjoint*, i.e., for every $i, j \in I$, $i \neq j$ implies $\llbracket [p_i, c_i] \rrbracket \cap \llbracket [p_j, c_j] \rrbracket = \emptyset$.
- (2) For every $\max X. \varphi$ we have $X \in \mathbf{fv}(\varphi)$. ■

In a (closed) SHML_{nf} formula, the basic terms tt and ff can never appear unguarded unless they are at the top level (e.g., we can never have $\varphi \wedge \text{ff}$ or $\max X_0. \dots \max X_n. \text{ff}$). Modal operators are combined with conjunctions into one construct $\bigwedge_{i \in I} [p_i, c_i] \varphi_i$ that is written as $[p_0, c_0] \varphi_0 \wedge \dots \wedge [p_n, c_n] \varphi_n$ when $I = \{0, \dots, n\}$ and simply as $[p_0, c_0] \varphi$ when $|I| = 1$. The conjunct modal guards must also be *disjoint* so that *at most one* necessity guard can satisfy any particular visible action. Along with these restrictions, we still assume that SHML_{nf} fixpoint variables are guarded, and that for every $((x)?(y), c)$, $y \notin \mathbf{fv}(c)$.

Example 5.2. The formula φ_3 defines a recursive property stating that an input on port \mathbf{a} (carrying any value) cannot be followed by an output with value of 4 (on any port), and that this continues to hold if the subsequent output is made on port \mathbf{a} with a value that is not equal to 3 (in which cases, the formula recurses)

$$\varphi_3 \stackrel{\text{def}}{=} \max X. \left[((x_1)?(y_1), x_1=\mathbf{a}) \left(\begin{array}{l} [((x_2)!(y_2), x_2=\mathbf{a} \wedge y_2 \neq 3)] X \\ \wedge [((x_3)!(y_3), y_3=4)] \text{ff} \end{array} \right) \right]$$

φ_3 is not an SHML_{nf} formula since its conjunction is not disjoint (e.g., the action $\mathbf{a}!4$ satisfies both branches). Still, we can reformulate φ_3 as $\varphi'_3 \in \text{SHML}_{\text{nf}}$:

$$\varphi'_3 \stackrel{\text{def}}{=} \max X. \left[((x_1)?(y_1), x_1=\mathbf{a}) \left(\begin{array}{l} [((x_4)!(y_4), x_4=\mathbf{a} \wedge y_4 \neq 4 \wedge y_4 \neq 3)] X \\ \wedge [((x_4)!(y_4), x_4=\mathbf{a} \wedge y_4=4)] \text{ff} \end{array} \right) \right]$$

where x_4 and y_4 are fresh variables. ■

Our monitor synthesis function in Definition 5.3 converts an SHML_{nf} formula φ into a transducer m . This conversion also requires information regarding the input ports employed by the SuS, as this is used to add the necessary insertion branches to silently unblock the SuS at runtime; this prevents the monitor from unnecessarily blocking the resulting composite system. . The synthesis function must therefore be supplied with this information in the form of a *finite* set of input ports $\Pi \subset \text{PORT}$, which then relays this information to the resulting monitor. It also assumes a default value v_{def} for the payload data domain.

²Recall that from Figure 2, I always denotes a *finite* set of indices which is crucial for a synthesis process to terminate.

Definition 5.3. The synthesis function $\llbracket - \rrbracket : \text{sHML}_{\text{nf}} \times \mathcal{P}_{\text{fin}}(\text{PORT}) \rightarrow \text{TRN}$ is defined inductively as:

$$\begin{aligned} \llbracket X, \Pi \rrbracket &\stackrel{\text{def}}{=} X \\ \llbracket \text{tt}, \Pi \rrbracket &\stackrel{\text{def}}{=} \text{id} \\ \llbracket \text{ff}, \Pi \rrbracket &\stackrel{\text{def}}{=} \text{id} \\ \llbracket \max X.\varphi, \Pi \rrbracket &\stackrel{\text{def}}{=} \text{rec } X.\llbracket \varphi, \Pi \rrbracket \\ \llbracket \varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rrbracket &\stackrel{\text{def}}{=} \text{rec } Y. \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i).\llbracket \varphi_i, \Pi \rrbracket & \text{otherwise} \end{cases} \right) + \text{def}(\varphi) \end{aligned}$$

$$\text{where } \text{dis}(p, c, m, \Pi) \stackrel{\text{def}}{=} \begin{cases} (p, c, \bullet).m & \text{if } p = (x)!(y) \\ \sum_{\mathbf{b} \in \Pi} (\bullet, c\{\mathbf{b}/x\}, \mathbf{b}?v_{\text{def}}).m & \text{if } p = (x)?(y) \end{cases}$$

and

$$\text{def}\left(\bigwedge_{i \in I} [(x_i)?(y_i), c_i]\varphi_i \wedge \psi\right) \stackrel{\text{def}}{=} \begin{cases} ((-)?(-)).\text{id} & \text{when } I = \emptyset \\ ((x)?(y), \bigwedge_{i \in I} (-c_i\{x/x_i, y/y_i\})).\text{id} & \text{otherwise} \end{cases}$$

where ψ has no conjuncts starting with an input modality, variables x and y are fresh, and v_{def} is a default value. \blacksquare

The definition above assumes a bijective mapping between formula variables and monitor recursion variables. Normalised conjunctions, $\bigwedge_{i \in I} [p_i, c_i]\varphi_i$, are synthesised as a *recursive summation* of monitors, i.e., $\text{rec } Y. \sum_{i \in I} m_i$, where Y is fresh, and every branch m_i can be one of the following:

- (i) when m_i is derived from a branch of the form $[p_i, c_i]\varphi_i$ where $\varphi_i \neq \text{ff}$, the synthesis produces a monitor with the *identity transformation* prefix, (p_i, c_i) , followed by the monitor synthesised from the continuation φ_i , i.e., $\llbracket \varphi_i, \Pi \rrbracket$;
- (ii) when m_i is derived from a violating branch of the form $[p_i, c_i]\text{ff}$, the synthesis produces an *action disabling transformation* via $\text{dis}(p_i, c_i, Y, \Pi)$.

Specifically, in clause (ii), the dis function produces either a *suppression transformation*, (p_i, c_i, \bullet) , when p_i is an *output* pattern, $(x_i)!(y_i)$, or a *summation of insertions*, $\sum_{\mathbf{b} \in \Pi} (\bullet, c_i\{\mathbf{b}/x_i\}, \mathbf{b}?v_{\text{def}}).m_i$, when p_i is an *input* pattern, $(x_i)?(y_i)$. The former signifies that the monitor must react to and suppress every matching (invalid) system output thus stopping it from reaching the environment. By not synthesising monitor branches that react to the erroneous input, the latter allows the monitor to hide the input synchronisations from the environment. At the same time, the synthesised insertion branches insert a default domain value v_{def} on every port $\mathbf{b} \in \Pi$ whenever the branch condition $c_i\{\mathbf{b}/x_i\}$ evaluates to true at runtime. This stops the monitor from blocking the runtime progression of the resulting composite system unnecessarily.

This blocking mechanism can, however, block *unspecified* inputs, i.e., those that do not satisfy any modal necessity in the normalised conjunction. This is undesirable since the unspecified actions do not contribute towards a safety violation and, instead, lead to its trivial satisfaction. To prevent this, the *default monitor* $\text{def}(\varphi)$ is also added to the resulting summation. Concretely, the def function produces a *catch-all* identity monitor that forwards an input to the SuS whenever it satisfies the negation of *all* the conditions associated with modal necessities for input patterns in the normalised conjunction. This condition is constructed for a normalised conjunction of the form $\bigwedge_{i \in I} [(x_i)?(y_i), c_i]\varphi_i \wedge \psi$

(assuming that ψ does not include further input modalities). Otherwise, if none of the conjunct modalities define an input pattern, every input is allowed, i.e., the default monitor becomes $((-)?(-)).\text{id}$, which transitions to id after forwarding the input to the SuS.

Example 5.4. Recall (the full version of) formula φ_1 from Example 2.4.

$$\begin{aligned}\varphi_1 &\stackrel{\text{def}}{=} \max X. [((x)?(y_1), x \neq \mathbf{b})] [(((x_1)?(-), x_1 = x)] \text{ff} \wedge [((x_2)!(y_2), x_2 = x)] \varphi'_1 \\ \varphi'_1 &\stackrel{\text{def}}{=} [((x_3)!(y_3), x_3 = x)] \text{ff} \wedge [((x_4)!(y_3), x_4 = \mathbf{b} \wedge y_3 = (\log, y_1, y_2))] X\end{aligned}$$

For any arbitrary set of ports Π , the synthesis of Definition 5.3 produces the following monitor.

$$\begin{aligned}m_{\varphi_1} &\stackrel{\text{def}}{=} \text{rec } X. \text{rec } Z. ((x)?(y_1), x \neq \mathbf{b}). \text{rec } Y_1. m'_{\varphi_1} + ((x_{\text{def}})?(-), x_{\text{def}} = \mathbf{b}). \text{id} \\ m'_{\varphi_1} &\stackrel{\text{def}}{=} \sum_{\mathbf{a} \in \Pi} (\bullet, \mathbf{a} = x, \mathbf{a} ? v_{\text{def}}). Y_1 + ((x_2)!(y_2), x_2 = x). \text{rec } Y_2. m''_{\varphi_1} + ((x_{\text{def}})?(-), x_{\text{def}} \neq x). \text{id} \\ m''_{\varphi_1} &\stackrel{\text{def}}{=} ((x_3)!(y_3), x_3 = x, \bullet). Y_2 + ((x_4)!(y_3), x_4 = \mathbf{b} \wedge y_3 = (\log, y_1, y_2)). X + ((-)?(-)). \text{id}\end{aligned}$$

Monitor m_{φ_1} can be optimised by removing redundant recursive constructs such as $\text{rec } Z.$ that are introduced mechanically by our synthesis. \blacksquare

Monitor m_{φ_1} from Example 5.4 (with $(\varphi_1, \Pi) = m_{\varphi_1}$) is very similar to m_{det} of Example 3.2, differing only in how it defines its insertion branches for unblocking the SuS. For instance, if we consider $\Pi = \{\mathbf{b}, \mathbf{c}\}$, (φ_1, Π) would synthesise two insertion branches, namely $(\bullet, \mathbf{b} = x, \mathbf{b} ? v_{\text{def}})$ and $(\bullet, \mathbf{c} = x, \mathbf{c} ? v_{\text{def}})$, but if Π also includes \mathbf{d} , it would add another branch. By contrast, the manually defined m_{det} attains the same result more succinctly via the single insertion branch $(\bullet, x ? v_{\text{def}})$. Importantly, our synthesis provides the witness monitors needed to show enforceability.

Theorem 5.5 (Enforceability). *sHML_{nf} is bidirectionally enforceable using the monitors and instrumentation of Figure 4.*

Proof. By Definition 4.1 the result follows from showing that for every $\varphi \in \text{sHML}_{\text{nf}}$ and $\Pi \subseteq \text{PORT}$, (φ, Π) enforces φ (for every Π). By Definition 4.9, this follows from Propositions 5.8 and 5.12, stated and proved in Section 5.1. \square

Theorem 5.5 entails that the synthesised monitors generated by the function described in Definition 5.3 do *enforce* their respective sHML_{nf} formula and are correct by construction. By this we mean that, if the formula φ being enforced can be expressed in the syntactic fragment sHML_{nf}, and it is satisfiable (i.e., $\llbracket \varphi \rrbracket \neq \emptyset$), then the resulting composite system, $m[s]$, consisting of the synthesised monitor, m , composed with the SuS, s , is guaranteed to satisfy φ and the changes to its original behaviour are only those that led to a violation. It is worth pointing out that should φ be unsatisfiable, there is very little that can be done by way of enforcement; the satisfiability caveats in Definitions 4.2, 4.4 and 4.6 are intentionally inserted so as not to require anything of the synthesised monitor in such cases. We argue that this way of dealing with unsatisfiable formulas is not a deficiency of the enforcement setup, but rather a flaw in the correctness specifications being imposed.

Note that the enforcement of formulas that use μHML constructs outside of the sHML is problematic. For instance, consider the disjunction formula $\varphi_1 \vee \varphi_2$ (recall that disjunctions are not part of the sHML syntax). In a branching-time setting, the subformulas φ_1 and φ_2 can, in principle, describe computation from different parts of the computation tree. This means that, although the current execution observed by a monitor might provide enough

$$\begin{aligned}
 (s, \text{tt}) \in \mathcal{R} & \text{ implies true} \\
 (s, \text{ff}) \in \mathcal{R} & \text{ implies false} \\
 (s, \bigwedge_{i \in I} \varphi_i) \in \mathcal{R} & \text{ implies } (s, \varphi_i) \in \mathcal{R} \text{ for all } i \in I \\
 (s, [(p, c)]\varphi) \in \mathcal{R} & \text{ implies } (\forall \alpha, r \cdot s \xrightarrow{\alpha} r \text{ and } (p, c)(\alpha) = \sigma) \text{ implies } (r, \varphi\sigma) \in \mathcal{R} \\
 (s, \max X.\varphi) \in \mathcal{R} & \text{ implies } (s, \varphi\{\max X.\varphi/X\}) \in \mathcal{R}
 \end{aligned}$$

where $(p, c)(\alpha) = \sigma$ is short for $\text{match}(p, \alpha) = \sigma$ and $c\sigma \Downarrow \text{true}$.

FIGURE 5. A satisfaction relation for sHML formulas

information to determine that one subformula is about to be violated (say φ_1), there could never be an execution that allows the monitor to determine when to intervene whenever *both* subformulas become violated. More precisely, by intervening to prevent φ_1 from being violated might break transparency, Definition 4.6, in cases where φ_2 is still satisfied (and thus $\varphi_1 \vee \varphi_2$ still holds). Conversely, not intervening might affect soundness, Definition 4.2, in cases where φ_2 is also violated (and thus $\varphi_1 \vee \varphi_2$ is certainly violated). It has been well established that a number of μ HML properties are not monitorable for a variety of settings [FAI17, AAFI18b, AAFI18a, AAF⁺19, AAF⁺21a, AAF⁺21b] and it is therefore reasonable to expect similar limits in the case of enforceability.

5.1. Enforceability Proofs. In what follows, we state and prove monitor soundness and transparency, Definitions 4.2 and 4.6 for the synthesis function presented in Definition 5.3. Upon first reading, the remainder of the section can be safely skipped without affecting the comprehension of the remaining material.

To facilitate the forthcoming proofs we occasionally use the satisfaction semantics for sHML from [AI99, HL95] which is defined in terms of the *satisfaction relation*, \models . When restricted to sHML, \models is the *largest relation* \mathcal{R} satisfying the implications defined in Figure 5. It is well known that this semantics agrees with the sHML semantics of Figure 2. As a result, we use $s \models \varphi$ in lieu of $s \in \llbracket \varphi \rrbracket$. At certain points in our proofs we also refer to the τ -closure property of sHML, Proposition 5.6, that was proven in [AI99].

Proposition 5.6. *if $s \xrightarrow{\tau} s'$ and $s \models \varphi$ then $s' \models \varphi$.* □

We start by stating and proving synthesis soundness, which relies on the following technical lemma relating recursive monitor unfolding and its behaviour.

Lemma 5.7. $\text{rec } X.m[s] \sim (m\{\text{rec } X.m/X\})[s]$

Proof. Follows from the instrumentation relation of Figure 4 and, more importantly, the monitor rule EREC, also in Figure 4. □

Proposition 5.8 (Soundness). *For every finite port set Π , system state $s \in \text{SYS}$ whose port names are included in Π , and $\varphi \in \text{sHML}_{nf}$, if $\llbracket \varphi \rrbracket \neq \emptyset$ then $(\varphi, \Pi)[s] \in \llbracket \varphi \rrbracket$.*

Proof. To prove that for every system s , formula φ and finite set of ports Π

$$\text{if } \llbracket \varphi \rrbracket \neq \emptyset \text{ then } (\varphi, \Pi)[s] \models \varphi$$

we setup the relation \mathcal{R} (below) and show that it is a *satisfaction relation* (\models) by demonstrating that it abides by the rules in Figure 5.

$$\mathcal{R} \stackrel{\text{def}}{=} \left\{ (r, \varphi) \left| \begin{array}{l} (i) \llbracket \varphi \rrbracket \neq \emptyset \text{ and } r = \langle \varphi, \Pi \rangle [s] \quad \text{or} \\ (ii) \llbracket \varphi \rrbracket \neq \emptyset \text{ and } r = \langle \max X_1 \dots \max X_n. \psi, \Pi \rangle [s] \\ \text{and } \varphi = \psi \{ \max X_1 \dots \max X_n. \psi / X_1 \} \dots \{ \max X_n. \psi / X_n \} \end{array} \right. \right\}$$

The second case defining the tuples $(r, \varphi) \in \mathcal{R}$ (labeled as *(ii)* for clarity) maps monitored system $m[s]$ where m is obtain by synthesising a formula consisting of a prefix of maximal fixpoint binders of length n , i.e., $m = \langle \varphi, \Pi \rangle$ where $\varphi = \max X_1 \dots \max X_n. \psi$, with the resp. *unfolded* formula $\psi \{ \max X_1 \dots \max X_n. \psi / X_1 \} \dots \{ \max X_n. \psi / X_n \}$.

We prove the claim that $\mathcal{R} \subseteq \models$ by case analysis on the structure of φ . We here consider the two main cases:

Case $\varphi = \max X. \psi$: We consider two subcases for why $(r, \varphi) \in \mathcal{R}$, following either condition *(i)* or *(ii)*:

Case *(i)*: We know that $\llbracket \max X. \psi \rrbracket \neq \emptyset$ and $r = \langle \max X. \psi, \Pi \rangle [s]$ for some s . By the rules defining (\models) in Figure 5 we need to show that

$$\langle \max X. \psi, \Pi \rangle [s], \psi \{ \max X. \psi / X \} \in \mathcal{R}$$

as well. This follows immediately from rule *(ii)* defining \mathcal{R} .

Case *(ii)*: We know that $\llbracket \max X. \psi \rrbracket \neq \emptyset$, that

$$\max X. \psi = \max X. (\phi \{ \max Y_1 \dots \max Y_k. \max X. \phi / Y_1 \} \dots \{ \max Y_k. \max X. \phi / Y_k \})$$

for some ϕ and k , and that $r = \langle \max Y_1 \dots \max Y_k. \max X. \phi, \Pi \rangle [s]$ for some s . Again, by the rules defining (\models) in Figure 5 we need to show that

$$\langle \max Y_1 \dots \max Y_k. \max X. \phi, \Pi \rangle [s], \phi' \in \mathcal{R}$$

for $\phi' = \phi \{ \max Y_1 \dots \max Y_k. \max X. \phi / Y_1 \} \dots \{ \max Y_k. \max X. \phi / Y_k \} \{ \max X. \phi / X \}$. This follows again from rule *(ii)* defining \mathcal{R} with an maximal fixpoint binder length set at $n = k + 1$.

Case $\varphi = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$ and $\#_{h \in I} (p_h, c_h)$: Again we have two subcases to consider for why we have the inclusion $(r, \varphi) \in \mathcal{R}$:

Case *(i)*: We know that

$$\llbracket \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i \rrbracket \neq \emptyset \tag{5.1}$$

and that $r = \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle [s]$ for some s . Recall that

$$\langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle = \text{rec } Y. \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i). \langle \varphi_i, \Pi \rangle & \text{otherwise} \end{cases} \right) + \text{def} \left(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i \right) \tag{5.2}$$

By the rules defining (\models) in Figure 5 (for the case involving $\bigwedge_{i \in I} \varphi_i$ and $[(p, c)] \varphi$ combined) we need to show that

$$\forall i \in I, \alpha, q \text{ if } \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle [s] \xrightarrow{\alpha} q \text{ and } (p_i, c_i)(\alpha) = \sigma \text{ then } (q, \varphi_i \sigma) \in \mathcal{R} \tag{5.3}$$

Pick any $[(p_i, c_i)] \varphi_i$ and proceed by case analysis:

Case $[(p_i, c_i)] \varphi_i = [(x)!(y), c_i] \text{ff}$: For any output action $a!v$ that the system s can produce, i.e., $s \xrightarrow{a!v} s'$, that matches the pattern of the necessity

formula considered, i.e., $((x)!(y), c_i)(a!v) = \sigma$, the monitor synthesised in Equation (5.2) transitions as

$$\langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle \xrightarrow{(a!v)\blacktriangleright\bullet} \langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle$$

Thus, by the instrumentation in Figure 4 (particularly rule `BIDISO`), we conclude that it could *never* be the case that

$$\langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle[s] \xrightarrow{a!v} q \text{ for any } q$$

meaning that condition (5.3) is satisfied.

Case $[(p_i, c_i)]\varphi_i = [(x)?(y), c_i]\text{ff}$: The reasoning is analogous to the previous case. For any input action $a?v$ that $s \xrightarrow{a?v} s'$, that matches the pattern of the necessity formula, $((x)?(y), c_i)(a?v) = \sigma$, the monitor synthesised in Equation (5.2) transitions as

$$\langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle \xrightarrow{\bullet(a?v)} \langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle$$

Thus, by the instrumentation in Figure 4 (particularly rule `BIDISI`), we conclude that it could *never* be the case that

$$\langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle[s] \xrightarrow{a?v} q \text{ for any } q$$

meaning that condition (5.3) is satisfied.

Case $\varphi_i \neq \text{ff}$: From Equation (5.1) we know that for any α such that $(p_i, c_i)(\alpha) = \sigma$ it holds that

$$\llbracket \varphi_i \sigma \rrbracket \neq \emptyset. \quad (5.4)$$

Now if $s \xrightarrow{\alpha} s'$, from the form of $\langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle$ in Equation (5.2) and $\#_{h \in I}(p_h, c_h)$ we conclude that

$$\langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle \xrightarrow{\alpha\blacktriangleright\alpha} \langle\langle \varphi_i, \Pi \rangle\rangle\sigma = \langle\langle \varphi_i \sigma, \Pi \rangle\rangle$$

Thus, by the instrumentation in Figure 4 (particularly rules `BITRNI` and `BITRNO`) we conclude

$$\langle\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle[s] \xrightarrow{\alpha} \langle\langle \varphi_i \sigma, \Pi \rangle\rangle[s']$$

and from Equation (5.4) and the definition of \mathcal{R} (i) we conclude that $(eBI(\langle\langle \varphi_i \sigma, \Pi \rangle\rangle s', \varphi_i \sigma) \in \mathcal{R}$, thus satisfying Equation (5.3) as required.

Case (ii): We know that $\llbracket \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i \rrbracket \neq \emptyset$, that for all $i \in I$

$$\varphi_i = \psi_i \{ \max Y_1 \dots \max Y_k. \psi_i / Y_1 \} \dots \{ \max Y_k. \psi_i / Y_k \} \text{ for some } \psi_i$$

and that $r = \langle\langle \max Y_1 \dots \max Y_k. \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle\rangle[s]$ for some s . Similar to the previous case, we need to show that r satisfies a requirement akin to Equation (5.3). This follows using a similar reasoning employed in the previous case, Lemma 5.7 and the transitivity of (strong) bisimulation. \square

We next state and prove the fact that the synthesis function of Definition 5.3 is eventually transparent, according to Definition 4.6. This proof for eventual transparency refers to the auxiliary Lemma 5.10 and another transparency result Proposition 5.11 (Transparency) for Definition 4.4, defined and proved below. The proof of Lemma 5.10, in turn, relies on the following technical lemma which states that any sequence τ transitions from a composite system enforced by a monitor synthesised from a conjuncted modal guard formula according to Definition 5.3 can be decomposed such that the monitored system is allowed to produce external actions by the monitor remains in the same state. which states that any sequence of τ transitions from a composite system enforced by a monitor synthesised from a conjunctive formula, according to Definition 5.3, stems from a corresponding sequence of external actions of the monitored system while the monitor remains in the same state.

Lemma 5.9. *For every formula of the form $\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$ and system states s and r , if $(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s] \xrightarrow{\tau} * r$ then there exists some state s' and trace u such that $s \xrightarrow{u} s'$ and $r = (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s']$.*

Proof. We proceed by mathematical induction on the number of τ transitions.

Case 0 transitions. This case holds trivially given that $s \xrightarrow{\varepsilon} s$ and so that $r = (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s]$.

Case $k + 1$ transitions. Assume that $(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s] \xrightarrow{\tau}^{k+1} r$ and so we can infer that

$$(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s] \xrightarrow{\tau} r' \quad (\text{for some } r') \quad (5.5)$$

$$r' \xrightarrow{\tau}^k r. \quad (5.6)$$

By the definition of $(-)$ we know that $(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)$ synthesises the monitor

$$\text{rec } Y. \sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i).(\varphi_i, \Pi) & \text{otherwise} \end{cases}$$

which can be unfolded into

$$(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi) = \sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, m, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i).(\varphi_i, \Pi) & \text{otherwise} \end{cases} \quad (5.7)$$

and so from (5.7) we know that the τ -reduction in (5.5) can be the result of rules IASY , IDISO or IDISI . We therefore inspect each case.

- IASY : By rule IASY , from (5.5) we can deduce that

$$\exists s'' . s \xrightarrow{\tau} s'' \quad (5.8)$$

$$r' = (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i)[s''] \quad (5.9)$$

and so by (5.6), (5.9) and the *inductive hypothesis* we know that

$$\exists s', u . s'' \xrightarrow{u} s' \text{ and } r = (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i)[s']. \quad (5.10)$$

Finally, by (5.8) and (5.10) we can thus conclude that $\exists s', u . s \xrightarrow{u} s'$ and also that $r = (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i)[s']$.

- **iDISI**: By rule **iDISI** and from (5.5) we infer that

$$\exists s'' \cdot s \xrightarrow{(a?v)} s'' \quad (5.11)$$

$$\langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle \xrightarrow{\bullet \blacktriangleright (a?v)} m' \quad (5.12)$$

$$r' = m'[s''] \quad (5.13)$$

and from (5.7) and by the definition of **dis** we can infer that the reduction in (5.12) occurs when the synthesised monitor inserts action $a?v$ and then reduces back to $\langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle$ allowing us to infer that

$$m' = \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle. \quad (5.14)$$

Hence, by (5.6), (5.13) and (5.14) we can apply the *inductive hypothesis* and deduce that

$$\exists s', u \cdot s'' \xrightarrow{u} s' \text{ and } r = \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle [s'] \quad (5.15)$$

so that by (5.11) and (5.15) we finally conclude that $\exists s', u \cdot s \xrightarrow{(a?v)u} s'$ and that $r = \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle [s']$ as required, and so we are done.

- **iDISO**: We omit the proof for this case as it is very similar to that of case **iDISI**. \square

The following lemma builds on Lemma 5.9, and states that the monitor obtained from a sequence of transitions t and a synthesised monitor $\langle \psi, \Pi \rangle$ can be calculated using the function $after(\varphi, t)$ and the synthesis function given in Definition 5.3.

Lemma 5.10. *For every set of names Π , formula $\varphi \in \text{SHML}_{\text{nf}}$, state s and trace t , if $\langle \varphi, \Pi \rangle [s] \xrightarrow{t} m'[s']$ then $\exists \psi \in \text{SHML}_{\text{nf}} \cdot \psi = after(\varphi, t)$ and $\langle \psi, \Pi \rangle = m'$.*

Proof. We need to prove that for every formula $\varphi \in \text{SHML}_{\text{nf}}$, if we assume that $\langle \varphi, \Pi \rangle [s] \xrightarrow{t} m'[s']$ then there must exist some formula ψ , such that $\psi = after(\varphi, t)$ and $\langle \psi, \Pi \rangle = m'$. We proceed by induction on the length of t .

Case $t = \varepsilon$. This case holds vacuously since when $t = \varepsilon$ then $m' = \langle \varphi, \Pi \rangle$ and $\varphi = after(\varphi, \varepsilon)$.

Case $t = \alpha u$. Assume that $\langle \varphi, \Pi \rangle [s] \xrightarrow{\alpha u} m'[s']$ from which by the definition \xrightarrow{t} we can infer that there are r and r' such that

$$\langle \varphi, \Pi \rangle [s] \xrightarrow{\tau} *r \quad (5.16)$$

$$r \xrightarrow{\alpha} r' \quad (5.17)$$

$$r' \xrightarrow{u} m'[s']. \quad (5.18)$$

We now proceed by case analysis on φ .

- $\varphi = X$: This case does not apply since $\langle \text{ff}, \Pi \rangle$ and $\langle X, \Pi \rangle$ do not yield a valid monitor.
- $\varphi \in \{\text{ff}, \text{tt}\}$: Since $\langle \text{tt}, \Pi \rangle = \text{id}$ we know that the τ -reductions in (5.16) are only possible via rule **iASY** which means that $s \xrightarrow{\tau} *s''$ and $r = \langle \text{tt}, \Pi \rangle [s'']$. The latter allows us to deduce that the reduction in (5.17) is only possible via rule **iTRN** and so we also know that

$s'' \xrightarrow{\alpha} *s'''$ and $r' = \llbracket \text{tt}, \Pi \rrbracket [s''']$. Hence, by (5.18) and the *inductive hypothesis* we conclude that

$$\exists \psi \in \text{sHML}_{\mathbf{nf}} \cdot \psi = \text{after}(\text{tt}, u) \quad (5.19)$$

$$\llbracket \psi, \Pi \rrbracket = m'. \quad (5.20)$$

Since from the definition of *after* we know that $\text{after}(\text{tt}, \alpha u)$ equates to $\text{after}(\text{after}(\text{tt}, \alpha), u)$ and $\text{after}(\text{tt}, \alpha) = \text{tt}$, from (5.19) we can conclude that $\psi = \text{after}(\text{tt}, \alpha u)$ and so this case holds since we also know (5.20). The case for **ff** is analogous.

- $\varphi = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$ and $\#_{i \in I} (p_i, c_i)$: Since $\varphi = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$, by the definition of $\llbracket - \rrbracket$ we know that $\text{rec } Y. \sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \mathbf{ff} \\ (p_i, c_i). \llbracket \varphi_i, \Pi \rrbracket & \text{otherwise} \end{cases}$ which can be unfolded into

$$\llbracket \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rrbracket = \sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, m, \Pi) & \text{if } \varphi_i = \mathbf{ff} \\ (p_i, c_i). \llbracket \varphi_i, \Pi \rrbracket & \text{otherwise} \end{cases} \quad (5.21)$$

and so by (5.16), (5.21) and Lemma 5.9 we conclude that $\exists s'' \cdot s \xrightarrow{u} s''$ and that

$$r = \llbracket \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rrbracket [s'']. \quad (5.22)$$

Hence, by (5.21) and (5.22) we know that the reduction in (5.17) can only happen if $\exists s'' \cdot s'' \xrightarrow{\alpha} s'''$ and α matches an identity transformation $(p_j, c_j). \llbracket \varphi_j, \Pi \rrbracket$ (for some $j \in I$) which was derived from $\llbracket (p_j, c_j) \rrbracket \varphi_j$ (where $\varphi_j \neq \mathbf{ff}$). We can thus deduce that

$$r' = \llbracket \varphi_j \sigma, \Pi \rrbracket [s'''] \quad (5.23)$$

$$\text{match}(p_j, \alpha) = \sigma \text{ and } c_j \sigma \Downarrow \text{true} \quad (5.24)$$

and so by (5.18), (5.23) and the *inductive hypothesis* we deduce that

$$\exists \psi \in \text{sHML}_{\mathbf{nf}} \cdot \psi = \text{after}(\varphi_j \sigma, u) \quad (5.25)$$

$$\llbracket \psi, \Pi \rrbracket = m'. \quad (5.26)$$

Now since we know (5.24), by the definition of *after* we infer that

$$\begin{aligned} \text{after}(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \alpha u) &= \text{after}(\text{after}(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \alpha), u) \\ &= \text{after}(\varphi_j \sigma, u) \end{aligned} \quad (5.27)$$

and so from (5.25) and (5.27) we conclude that

$$\psi = \text{after}(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \alpha u). \quad (5.28)$$

Hence, this case is done by (5.26) and (5.28).

- $\varphi = \max X. \psi$ and $X \in \mathbf{fv}(\psi)$: Since $\varphi = \max X. \psi$, by the syntactic rules of $\text{sHML}_{\mathbf{nf}}$ we know that $\psi \notin \{\mathbf{ff}, \text{tt}\}$ since $X \notin \mathbf{fv}(\psi)$, and that $\psi \neq X$ since logical variables must be guarded, hence we know that ψ can only be of the form

$$\psi = \max Y_1 \dots \max Y_n. \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i. \quad (5.29)$$

where $\max Y_1 \dots \max Y_n$ denotes an arbitrary number of fixpoint declarations, possibly none. Hence, knowing (5.29), by unfolding every fixpoint in $\max X. \psi$ we reduce the

formula to

$$\varphi = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i \{ \max X. \max Y_1 \dots \max Y_n. \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i / X, \dots \}$$

and so from this point onwards the proof proceeds as per that of case $\varphi = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$ which allows us to deduce that

$$\exists \psi' \in \text{SHML}_{\text{nf}} \cdot \psi' = \text{after}(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i \{ \dots \}, \alpha u) \quad (5.30)$$

$$\langle \psi', \Pi \rangle = m'. \quad (5.31)$$

From (5.29), (5.30) and the definition of *after* we can therefore conclude that

$$\exists \psi' \in \text{SHML}_{\text{nf}} \cdot \psi' = \text{after}(\max X. \psi, \alpha u) \quad (5.32)$$

and so this case holds by (5.31) and (5.32).

Hence, the above cases suffice to show that the case for when $t = \alpha u$ holds. \square

The transparency proof following Definition 4.4 is given below.

Proposition 5.11 (Transparency). *For every state $s \in \text{SYS}$ and $\varphi \in \text{SHML}_{\text{nf}}$, if $s \in \llbracket \varphi \rrbracket$ then $\langle \varphi, \Pi \rangle [s] \sim s$.*

Proof. Since $s \in \llbracket \varphi \rrbracket$ is analogous to $s \models \varphi$ we prove that relation $\mathcal{R} \stackrel{\text{def}}{=} \{ (s, \langle \varphi, \Pi \rangle [s]) \mid s \models \varphi \}$ is a *strong bisimulation relation* that satisfies the following transfer properties:

- (a) if $s \xrightarrow{\mu} s'$ then $\langle \varphi, \Pi \rangle [s] \xrightarrow{\mu} r'$ and $(s', r') \in \mathcal{R}$
- (b) if $\langle \varphi, \Pi \rangle [s] \xrightarrow{\mu} r'$ then $s \xrightarrow{\mu} s'$ and $(s', r') \in \mathcal{R}$

We prove (a) and (b) separately by assuming that $s \models \varphi$ in both cases as defined by relation \mathcal{R} and conduct these proofs by case analysis on φ . We now proceed to prove (a) by case analysis on φ .

Cases $\varphi \in \{\text{ff}, X\}$. Both cases do not apply since $\nexists s \cdot s \models \text{ff}$ and similarly since X is an open-formula and so $\nexists s \cdot s \models X$.

Case $\varphi = \text{tt}$. We now assume that

$$s \models \text{tt} \quad (5.33)$$

$$s \xrightarrow{\mu} s' \quad (5.34)$$

and since $\mu \in \{\tau, \alpha\}$, we must consider both cases.

- $\mu = \tau$: Since $\mu = \tau$, we can apply rule IASY on (5.34) and get that

$$\langle \text{tt}, \Pi \rangle [s] \xrightarrow{\tau} \langle \text{tt}, \Pi \rangle [s'] \quad (5.35)$$

as required. Also, since we know that every process satisfies tt , we know that $s' \models \text{tt}$, and so by the definition of \mathcal{R} we conclude that

$$(s', \langle \text{tt}, \Pi \rangle [s']) \in \mathcal{R} \quad (5.36)$$

as required. This means that this case is done by (5.35) and (5.36).

- $\mu = \alpha$: Since $\langle \text{tt}, \Pi \rangle = \text{id}$ encodes the ‘catch-all’ monitor, $\text{rec } Y. ((x)! (y), \text{true}, x! y). Y + ((x)? (y), \text{true}, x? y). Y$, by rules EREC and ETRN we can apply rule ITRNI/O and deduce

that $\text{id} \xrightarrow{\alpha \blacktriangleright \alpha} \text{id}$, which we can further refine as

$$\langle \text{tt}, \Pi \rangle [s] \xrightarrow{\alpha} \langle \text{tt}, \Pi \rangle [s'] \quad (5.37)$$

as required. Once again since $s' \models \text{tt}$, by the definition of \mathcal{R} we can infer that

$$(s', \langle \text{tt}, \Pi \rangle [s']) \in \mathcal{R} \quad (5.38)$$

as required, and so this case is done by (5.37) and (5.38).

Case $\varphi = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$. We assume that

$$s \models \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i \quad (5.39)$$

$$s \xrightarrow{\mu} s' \quad (5.40)$$

and by the definition of \models and (5.39) we have that for every index $i \in I$ and action $\beta \in \text{ACT}$,

$$\text{if } s \xrightarrow{\beta} s' \text{ and } (p_i, c_i)(\beta) = \sigma \text{ then } s' \models \varphi_i \sigma. \quad (5.41)$$

Since $\mu \in \{\tau, \alpha\}$, we must consider both possibilities.

- $\mu = \tau$: Since $\mu = \tau$, we can apply rule IASY on (5.40) and obtain

$$\langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle [s] \xrightarrow{\tau} \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle [s''] \quad (5.42)$$

as required. Since $\mu = \tau$, and since we know that sHML is τ -closed, from (5.39), (5.40) and Proposition 5.6, we can deduce that $s' \models \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$, so that by the definition of \mathcal{R} we conclude that

$$(s'', \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle [s'']) \in \mathcal{R} \quad (5.43)$$

as required. This subcase is therefore done by (5.42) and (5.43).

- $\mu = \alpha$: Since $\mu = \alpha$, from (5.40) we know that

$$s \xrightarrow{\alpha} s' \quad (5.44)$$

and by the definition of $\langle - \rangle$ we can immediately deduce that

$$\langle \varphi_{\wedge}, \Pi \rangle = \text{rec } Y. \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i). \langle \varphi_i, \Pi \rangle & \text{otherwise} \end{cases} \right) + \text{def}(\varphi_{\wedge}) \quad (5.45)$$

where $\varphi_{\wedge} \stackrel{\text{def}}{=} \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$. Since the branches in the conjunction are all disjoint, $\#_{i \in I} (p_i, c_i)$, we know that *at most one* of the branches can match the same (input or output) action α . Hence, we consider two cases, namely:

- *No matching branches* (i.e., $\forall i \in I \cdot (p_i, c_i)(\alpha) = \text{undef}$): Since none of the symbolic actions in (5.45) can match action α , we can infer that if α is an *input*, i.e., $\alpha = \mathbf{a}!v$, then it will match the default monitor $\text{def}(\varphi_{\wedge})$ and transition via rule ITRNI, while if it is an *output*, i.e., $\alpha = \mathbf{a}!v$, rule IDEF handles the underspecification. In both cases, the monitor reduces to id . Also, notice that rules IDISO and IDISI cannot be applied since if they do, it would mean that s can also perform an erroneous action, which is not the

case since we assume (5.39). Hence, we infer that

$$\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle [s] \xrightarrow{\alpha} \langle \text{tt}, \Pi \rangle [s'] \quad (\text{since id} = \langle \text{tt}, \Pi \rangle) \quad (5.46)$$

as required. Also, since any process satisfies tt , we know that $s' \models \text{tt}$, and so by the definition of \mathcal{R} we conclude that

$$(s', \langle \text{tt}, \Pi \rangle [s']) \in \mathcal{R} \quad (5.47)$$

as required. This case is therefore done by (5.46) and (5.47).

- *One matching branch (i.e., $\exists j \in I \cdot (p_j, c_j)(\alpha) = \sigma$):* From (5.45) we can infer that the synthesised monitor can only disable the (input or output) actions that are defined by violating modal necessities. However, from (5.41) we also deduce that s is *incapable* of executing such an action as that would contradict assumption (5.39). Hence, since we now assume that $\exists j \in I \cdot (p_j, c_j)(\alpha) = \sigma$, from (5.45) we deduce that this action can only be transformed by an identity transformation and so by rule ETrn we have that

$$(p_j, c_j) \cdot \langle \varphi_j, \Pi \rangle \xrightarrow{\alpha \blacktriangleright \alpha} \langle \varphi_j \sigma, \Pi \rangle. \quad (5.48)$$

By applying rules ESEL , EREC on (5.48) and by (5.44), (5.45) and ITrNI/O (depending on whether α is an input or output action) we get that

$$\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle [s] \xrightarrow{\alpha} \langle \varphi_j \sigma, \Pi \rangle [s'] \quad (5.49)$$

as required. By (5.41), (5.44) and since we assume that $\exists j \in I \cdot (p_j, c_j)(\alpha) = \sigma$ we have that $s' \models \varphi_j \sigma$, and so by the definition of \mathcal{R} we conclude that

$$(s', \langle \varphi_j \sigma, \Pi \rangle [s']) \in \mathcal{R} \quad (5.50)$$

as required. Hence, this subcase holds by (5.49) and (5.50).

Case $\varphi = \max X.\varphi$ and $X \in \mathbf{fv}(\varphi)$. Now, let's assume that

$$s \xrightarrow{\mu} s' \quad (5.51)$$

and that $s \models \max X.\varphi$ from which by the definition of \models we have that

$$s \models \varphi\{\max X.\varphi/X\}. \quad (5.52)$$

Since $\varphi\{\max X.\varphi/X\} \in \text{SHML}_{\text{nf}}$, by the restrictions imposed by SHML_{nf} we know that: φ cannot be X because (bound) logical variables are required to be *guarded*, and it also cannot be tt or ff since X is required to be defined in φ , i.e., $X \in \mathbf{fv}(\varphi)$. Hence, we know that φ can only have the following form, that is

$$\varphi = \max Y_0. \dots \max Y_n. \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i \quad (5.53)$$

and so by (5.52), (5.53) and the definition of \models we have that

$$s \models (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i)\{\dots\} \quad \text{where} \quad (5.54)$$

$$\{\dots\} = \{\max X.\varphi/X, (\max Y_0. \dots \max Y_n. \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i)/Y_0, \dots\}.$$

Since we know (5.51) and (5.54), from this point onwards the proof proceeds as in the previous case. We thus omit the details.

These cases thus allow us to conclude that (a) holds. We now proceed to prove (b) using a similar case analysis approach.

Cases $\varphi \in \{\text{ff}, X\}$. Both cases do not apply since $\nexists s \cdot s \models \text{ff}$ and similarly since X is an open-formula and $\nexists s \cdot s \models X$.

Case $\varphi = \text{tt}$. Assume that

$$s \models \text{tt} \tag{5.55}$$

$$\langle \text{tt}, \Pi \rangle [s] \xrightarrow{\mu} r' \tag{5.56}$$

Since $\mu \in \{\tau, a?v, a!v\}$, we must consider each case.

- $\mu = \tau$: Since $\mu = \tau$, the transition in (5.56) can be performed via iDISI , iDISO or iASY . We must therefore consider these cases.

- iASY : From rule iASY and (5.56) we thus know that $r' = \langle \text{tt}, \Pi \rangle [s']$ and that $s \xrightarrow{\tau} s'$ as required. Also, since every process satisfies tt , we know that $s' \models \text{tt}$ as well, and so we are done since by the definition of \mathcal{R} we know that $(s', \langle \text{tt}, \Pi \rangle [s']) \in \mathcal{R}$.
- iDISI : From rule iDISI and (5.56) we know that: $r' = m'[s']$, $s \xrightarrow{a?v} s'$ and that

$$\langle \text{tt}, \Pi \rangle \xrightarrow{\bullet \blacktriangleright (a?v)} m'. \tag{5.57}$$

Since $\langle \text{tt}, \Pi \rangle = \text{id}$ we can deduce that (5.57) is *false* and hence this case does not apply.

- iDISO : The proof for this case is analogous as to that of case iDISI .

- $\mu = a?v$: Since $\mu = a?v$, the transition in (5.56) can be performed either via iTRNI or iENI . We consider both cases.

- iENI : This case also does not apply since if the transition in (5.56) is caused by rule iENI we would have that $\langle \text{tt}, \Pi \rangle \xrightarrow{a?v \blacktriangleright \bullet} m$ which is *false* since $\langle \text{tt}, \Pi \rangle = \text{id} = \text{rec } Y.((x)!(y), \text{true}, x!y).Y + ((x)?(y), \text{true}, x?y).Y$ and rules EREC , ESEL and ETRN state that for every $a?v$, $\text{id} \xrightarrow{a?v \blacktriangleright a?v} \text{id}$, thus leading to a contradiction.
- iTRNI : By applying rule iTRNI on (5.56) we know that $r' = m'[s']$ such that

$$\langle \text{tt}, \Pi \rangle \xrightarrow{a?v \blacktriangleright b?w} m'. \tag{5.58}$$

$$s \xrightarrow{b?w} s' \tag{5.59}$$

Since $\langle \text{tt}, \Pi \rangle = \text{id} = \text{rec } Y.((x)!(y), \text{true}, x!y).Y + ((x)?(y), \text{true}, x?y).Y$, by applying rules EREC , ESEL and ETRN to (5.58) we know that $a?v = b?w$, $m' = \text{id} = \langle \text{tt}, \Pi \rangle$, meaning that $r' = \langle \text{tt}, \Pi \rangle [s']$. Hence, since every process satisfies tt we know that $s' \models \text{tt}$, so that by the definition of \mathcal{R} we conclude

$$(s', \langle \text{tt}, \Pi \rangle [s']) \in \mathcal{R}. \tag{5.60}$$

Hence, we are done by (5.59) and (5.60) since we know that $a?v = b?w$.

- $\mu = a!v$: When $\mu = a!v$, the transition in (5.56) can be performed via iDEF , iTRNO or iENO . We omit this proof as it is very similar to that of case $\mu = a?v$.

Case $\varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$. We now assume that

$$s \models \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i \quad (5.61)$$

$$(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s] \xrightarrow{\mu} r'. \quad (5.62)$$

From (5.61) and by the definition of \models we can deduce that

$$\forall i \in I, \beta \in \text{ACT} \cdot \text{if } s \xrightarrow{\beta} s' \text{ and } (p_i, c_i)(\beta) = \sigma \text{ then } s' \models \varphi_i \sigma \quad (5.63)$$

and from (5.62) and the definition of $(-)$ we have that

$$\left(\text{rec } Y. \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i).(\varphi_i, \Pi) & \text{otherwise} \end{cases} \right) + \text{def}(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i) \right) [s'] \xrightarrow{\mu} r'. \quad (5.64)$$

From (5.64) we can deduce that the synthesised monitor can only disable an (input or output) action β when its occurrence would violate a conjunct of the form $[(p_i, c_i)]\text{ff}$ for some $i \in I$. However, we also know that s is *unable* to perform such an action as otherwise it would contradict assumption (5.63). Hence, we can safely conclude that the synthesised monitor in (5.64) does *not* disable any (input or output) actions of s , and so by the definition of dis we conclude that

$$\forall a?v, a!v \in \text{ACT}, s' \in \text{SYS} \cdot \left(\begin{array}{l} s \xrightarrow{a?v} s' \text{ implies } (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi) \xrightarrow{\bullet a?w} \text{ (for all } w \text{) and} \\ s \xrightarrow{a!v} s' \text{ implies } (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi) \xrightarrow{a?v \bullet} \end{array} \right). \quad (5.65)$$

Since $\mu \in \{\tau, a?v, a!v\}$, we must consider each case.

- $\mu = \tau$: Since $\mu = \tau$, from (5.62) we know that

$$(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s] \xrightarrow{\tau} r' \quad (5.66)$$

The τ -transition in (5.66) can be the result of rules iASY , iDISI or iDISO ; we thus consider each eventuality.

- iASY : As we assume that the reduction in (5.66) is the result of rule iASY , we know that $r' = (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s']$ and that

$$s \xrightarrow{\tau} s' \quad (5.67)$$

as required. Also, since sHML is τ -closed, by (5.61), (5.67) and Proposition 5.6 we deduce that $s' \models \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$ as well, so that by the definition of \mathcal{R} we conclude that

$$(s', (\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi)[s']) \in \mathcal{R} \quad (5.68)$$

and so we are done by (5.67) and (5.68).

- iDISI : By assuming that reduction (5.66) results from iDISI , we have that $r' = m'[s']$ and that

$$(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi) \xrightarrow{\bullet a?v} m' \quad (5.69)$$

$$s \xrightarrow{a?v} s' \quad (5.70)$$

By (5.65) and (5.70) we can, however, deduce that for every value w , we have that

$(\llbracket \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rrbracket \xrightarrow{\bullet \blacktriangleright a^?w}$. This contradicts (5.69) and so this case does not apply.

- **iDISO**: As we now assume that the reduction in (5.66) results from **iDISO**, we have that $r' = m'[s']$ and that

$$s \xrightarrow{a!v} s' \quad (5.71)$$

$$(\llbracket \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rrbracket \xrightarrow{a!v \blacktriangleright \bullet} m'. \quad (5.72)$$

Again, this case does not apply since from (5.65) and (5.71) we can deduce that

$(\llbracket \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rrbracket \xrightarrow{a!v \blacktriangleright \bullet}$ which contradicts (5.72).

- $\mu = a^?v$: When $\mu = a^?v$, the transition in (5.64) can be performed via rules **iENI** or **iTRNI**, we consider both possibilities.
 - **iENI**: This case does not apply since from (5.64) and by the definition of $(\llbracket - \rrbracket)$ we know that the synthesised monitor does not include action enabling transformations.
 - **iTRNI**: By assuming that (5.64) is obtained from rule **iTRNI** we know that

$$\text{rec } Y. \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i).(\llbracket \varphi_i, \Pi \rrbracket) & \text{otherwise} \end{cases} \right) + \text{def}(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i) \xrightarrow{a^?v \blacktriangleright b^?w} m' \quad (5.73)$$

$$s \xrightarrow{b^?w} s' \quad (5.74)$$

$$r' = m'[s']. \quad (5.75)$$

Since from (5.65) we know that the synthesised monitor in (5.73) does not disable any action performable by s , and since from the definition of $(\llbracket - \rrbracket)$ we know that the synthesis is incapable of producing action replacing monitors, we can deduce that

$$a^?v = b^?w. \quad (5.76)$$

With the knowledge of (5.76), from (5.74) we can thus deduce that

$$s \xrightarrow{a^?v} s' \quad (5.77)$$

as required. Knowing (5.76) we can also deduce that in (5.73) the monitor transforms an action $a^?v$ either (i) via an identity transformation that was synthesised from one of the *disjoint* conjunction branches, *i.e.*, from a branch $(p_j, c_j).(\llbracket \varphi_j, \Pi \rrbracket)$ for some $j \in I$, or else (ii) via the default monitor synthesised by $\text{def}(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i)$. We consider both eventualities.

- (i) In this case we apply rules **EREC**, **ESEL** and **ETRNI** on (5.73) and deduce that

$$\exists j \in I. (p_j, c_j)(a^?v) = \sigma \quad (5.78)$$

$$m' = (\llbracket \varphi_j \sigma, \Pi \rrbracket). \quad (5.79)$$

and so from (5.77), (5.78) and (5.63) we infer that $s' \models \varphi_j \sigma$ from which by the definition of \mathcal{R} we have that $(s', (\llbracket \varphi_j \sigma, \Pi \rrbracket)[s']) \in \mathcal{R}$, and so from (5.75) and (5.79) we can conclude that

$$(s', r') \in \mathcal{R} \quad (5.80)$$

as required, and so this case is done by (5.77) and (5.80).

- (ii) When we apply rules EREC, ESEL and ETRN we deduce that $m' = \text{id}$ and so by the definition of $\langle - \rangle$ we have that

$$m' = \langle \text{tt}, \Pi \rangle. \quad (5.81)$$

Consequently, as every process satisfies tt , we know that $s' \models \text{tt}$ and so by the definition of \mathcal{R} we have that $(s', \langle \text{tt}, \Pi \rangle[s']) \in \mathcal{R}$, so that from (5.75) and (5.81) we can conclude that

$$(s', r') \in \mathcal{R} \quad (5.82)$$

as required. Hence this case is done by (5.77) and (5.82).

- $\mu = a!v$: When $\mu = a!v$, the transition in (5.64) can be performed via IDEF, ITRNO or IENO. We omit the proof for this case due to its strong resemblance to that of case $\mu = a?v$.

Case $\varphi = \max X.\varphi$ and $X \in \mathbf{fv}(\varphi)$. Now, let's assume that

$$\langle \max X.\varphi, \Pi \rangle[s] \xrightarrow{\mu} r' \quad (5.83)$$

and that $s \models \max X.\varphi$ from which by the definition of \models we have that

$$s \models \varphi\{\max X.\varphi/X\}. \quad (5.84)$$

Since $\varphi\{\max X.\varphi/X\} \in \text{SHML}_{\mathbf{nf}}$, by the restrictions imposed by $\text{SHML}_{\mathbf{nf}}$ we know that: φ cannot be X because (bound) logical variables are required to be *guarded*, and it also cannot be tt or ff since X is required to be defined in φ , i.e., $X \in \mathbf{fv}(\varphi)$. Hence, we know that φ can only have the following form, that is

$$\varphi = \max Y_0. \dots \max Y_n. \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i \quad (5.85)$$

and so by (5.84), (5.85) and the definition of \models we have that

$$s \models \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i\{\dots\} \quad \text{where} \quad (5.86)$$

$$\{\dots\} = \{\max X.\varphi/X, (\max Y_0. \dots \max Y_n. \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i)/Y_0, \dots\}.$$

Since $\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i\{\dots\}, \Pi \rangle$ synthesises the *unfolded equivalent* of $\langle \max X.\varphi, \Pi \rangle$, from (5.83) we know that

$$\langle \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i\{\dots\}, \Pi \rangle[s] \xrightarrow{\mu} r'. \quad (5.87)$$

Hence, since we know (5.86) and (5.87), from this point onwards the proof proceeds as per the previous case. We thus omit showing the remainder of this proof.

From the above cases we can therefore conclude that (b) holds as well. \square

We are finally in a position to state and prove our eventual transparency results following Definition 4.6.

$$mc(m, t_\tau) \stackrel{\text{def}}{=} \begin{cases} 1 + mc(m', t'_\tau) & \text{if } t_\tau = \mu t'_\tau \text{ and } m[\text{sys}(\mu t'_\tau)] \xrightarrow{\mu'} m'[\text{sys}(t'_\tau)] \text{ and } \mu \neq \mu' \\ 1 + mc(m', t_\tau) & \text{if } t_\tau \in \{\mu t'_\tau, \varepsilon\} \text{ and } m[\text{sys}(t_\tau)] \xrightarrow{\mu'} m'[\text{sys}(t_\tau)] \\ mc(m', t'_\tau) & \text{if } t_\tau = \mu t'_\tau \text{ and } m[\text{sys}(\mu t'_\tau)] \xrightarrow{\mu} m'[\text{sys}(t'_\tau)] \\ |t_\tau| & \text{if } t_\tau \in \{\mu t'_\tau, \varepsilon\} \text{ and } \forall \mu' \cdot m[\text{sys}(t_\tau)] \not\xrightarrow{\mu'} \end{cases}$$

FIGURE 6. Modification Count (mc).

Proposition 5.12 (Eventual Transparency). *For every input port set Π , sHML_{nf} formula φ , system states $s, s' \in \text{Sys}$, action disabling monitor m' and trace t , if $\langle \varphi, \Pi \rangle[s] \xrightarrow{t} m'[s']$ and $s' \in \llbracket \text{after}(\varphi, t) \rrbracket$ then $m'[s'] \sim s'$. \square*

Proof. We must prove that for every formula $\varphi \in \text{sHML}_{\text{nf}}$ if $\langle \varphi, \Pi \rangle = m$ then $\text{evtenf}(m, \varphi)$. We prove that for every $\varphi \in \text{sHML}_{\text{nf}}$, if $\langle \varphi, \Pi \rangle[s] \xrightarrow{t} m'[s']$ and $s' \models \text{after}(\varphi, t)$ then $m'[s'] \sim s'$.

Now, assume that

$$\langle \varphi, \Pi \rangle[s] \xrightarrow{t} m'[s'] \tag{5.88}$$

$$s' \models \text{after}(\varphi, t) \tag{5.89}$$

and so from (5.88) and Lemma 5.10 we have that

$$\exists \psi \in \text{sHML}_{\text{nf}} \cdot \psi = \text{after}(\varphi, t) \tag{5.90}$$

$$\langle \text{after}(\varphi, t), \Pi \rangle = m' = \langle \psi, \Pi \rangle. \tag{5.91}$$

Hence, knowing (5.89) and (5.90), by Proposition 5.11 (Transparency) we conclude that $\langle \text{after}(\varphi, t), \Pi \rangle[s'] \sim s'$ as required, and so we are done. \square

6. TRANSDUCER OPTIMALITY

Recall Definition 4.9 from Section 4. Through criteria such as Definitions 4.2 and 4.6, it defined what it means for a monitor to adequately enforce a formula. However, it did *not* assess whether a monitor is (to some extent) the “*best*” that one can find to enforce a property. In order to define such a notion we must first be able to compare monitors to one another via some kind of *distance measurement* that tells them apart. One potential measurement is to assess the monitor’s level of *intrusiveness* when enforcing the property.

In Figure 6 we define function mc that inductively analyses a system run, represented as an explicit trace t_τ , and counts the number of modifications applied by the monitor. In each case the function reconstructs a trace system $\text{sys}(t_\tau)$ and instruments it with the monitor m in order to assess the type of transformation applied. Specifically, in the first two cases, mc increments the counter whenever the monitor adapts, disables or enables an action, and then it recurses to keep on inspecting the run (*i.e.*, the suffix t'_τ in the first, and the same trace t_τ in the second) vis-a-vis the subsequent monitor state, m' . The third case, specifies that the counter stays unmodified when the monitor applies an identity transformation, while the last case returns the length of t_τ when $m[\text{sys}(t_\tau)]$ is unable to execute further.

$$ec(m) \stackrel{\text{def}}{=} \begin{cases} \emptyset & \text{if } m = X \\ \bigcup_{i \in I} ec(m_i) & \text{if } m = \sum_{i \in I} m_i \\ ec(m') & \text{if } m = \text{rec } X.m' \text{ or } m = (p, c, \underline{p}).m' \\ \{\text{DIS}\} \cup ec(m') & \text{if } m = ((x)!(y), c, \bullet).m' \text{ or } m = (\bullet, c, a?v).m' \\ \{\text{EN}\} \cup ec(m') & \text{if } m = ((x)?(y), c, \bullet).m' \text{ or } m = (\bullet, c, a!v).m' \\ \{\text{ADPT}\} \cup ec(m') & \text{if } m = (p, c, p').m' \text{ and } p' \neq \underline{p} \neq \bullet \end{cases}$$

 FIGURE 7. Enforcement Capabilities (ec).

Example 6.1. Recall the monitors of Example 3.2 and consider the following system run $t_\tau^0 = a?v_1.a?v_2.\tau.a!w_2.a!w_2.b!(\log, v_2, w_2)$. For m_e and m_a , function mc respectively counts three enabled actions, i.e., $mc(m_e, t_\tau^0) = 3$, and four adapted actions, i.e., $mc(m_a, t_\tau^0) = 4$ (since $b!(\log, v_2, w_2)$ remains unmodified). The maximum count of 5 is attained by m_d as it immediately blocks the first input $a?v_1$, and so none of the actions in t_τ^0 can be executed by the composite system i.e., $\forall \mu \cdot m_d[\text{sys}(t_\tau^0)] \not\rightarrow \mu$ and so $mc(m_d, t_\tau^0) = 5$. Similarly, $mc(m_{dt}, t_\tau^0) = 4$ since m_{dt} allows the first request to be made, but blocks the second erroneous one, and as a result it also forbids the execution of the subsequent actions, i.e., $\forall \mu \cdot m_{dt}[\text{sys}(t_\tau^0)] \xrightarrow{a?v_1} \cdot \not\rightarrow \mu$. Finally, m_{det} performs the least number of modifications, namely $mc(m_{det}, t_\tau^0) = 2$. The first modification is caused when the monitor blocks the second erroneous input and internally *inserts* a default input value that allows the composite system to proceed over a τ -action. This contrasts with m_d and m_{dt} which fail to perform this insertion step thereby contributing to their high intrusiveness score. The second modification is attained when m_{det} suppresses the redundant response. ■

We can now use function mc to compare monitors to each other in order to identify the least intrusive one, i.e., the monitor that applies the least amount of transformations when enforcing a specific property. However, for this comparison to be fair, we must also compare like with like. This means that if a monitor enforces a formula by only disabling actions, it is only fair to compare it to other monitors of the same kind. It is reasonable to expect that monitors with more enforcement capabilities are likely to be “better” than those with fewer capabilities. We determine the enforcement capabilities of a monitor via function ec of Figure 7. It inductively analyses the structure of the monitor and deduces whether it can enable, disable and adapt actions based on the type of transformation triples it defines. For instance, if the monitor defines an output suppression triple, $((x)!(y), c, \bullet).m'$, or an input insertion branch, $(\bullet, c, a?v).m'$, then ec determines that the monitor can disable actions DIS, while if it defines an input suppression, $((x)?(y), c, \bullet).m'$, or an output insertion branch, $(\bullet, c, a!v).m'$, then it concludes that the monitor can enable actions, EN. Similarly, if a monitor defines a replacement transformation, it infers that the monitor can adapt actions, ADPT.

Example 6.2. Recall the monitors of Example 3.2. With function ec we determine that $ec(m_e) = \{\text{EN}\}$, $ec(m_a) = \{\text{ADPT}\}$, $ec(m_d) = ec(m_{dt}) = ec(m_{det}) = \{\text{DIS}\}$. Monitors may also have multiple types of enforcement capabilities. For instance,

$$ec(\text{rec } X.((x)?(y), \bullet).X + ((x)!(y), \bullet).X) = \{\text{EN}, \text{DIS}\}. \quad \blacksquare$$

With these definitions we now define *optimal enforcement*.

Definition 6.3 (Optimal Enforcement). A monitor m is *optimal* when enforcing φ , denoted as $\text{oenf}(m, \varphi)$, iff it *enforces* φ (Definition 4.9) and when for every state s , explicit trace t_τ and monitor n , if $ec(n) \subseteq ec(m)$, $\text{enf}(n, \varphi)$ and $s \xrightarrow{t_\tau}$ then $mc(m, t_\tau) \leq mc(n, t_\tau)$. \square

Definition 6.3 states that an adequate (sound and eventually transparent) monitor m is *optimal* for φ , if one cannot find another adequate monitor n , with the same (or fewer) enforcement capabilities, that performs *fewer modifications* than m and is thus *less intrusive*.

Example 6.4. Recall formula φ_1 of Example 2.4 and monitor m_{det} of Example 4.7. Although showing that $\text{oenf}(m_{\text{det}}, \varphi_1)$ is inherently difficult, from Example 6.1 we already get the intuition that it holds since m_{det} imposes the least amount of modifications compared to the other monitors of Examples 3.2 and 3.3. We further reaffirm this intuition using systems $s_{\mathbf{b}}$ and $s_{\mathbf{c}}$ from Example 2.4. In fact, when considering the invalid runs $t_\tau^1 \stackrel{\text{def}}{=} \mathbf{a}?v_1.\tau.\mathbf{a}!w_1.\mathbf{a}!w_1.\mathbf{b}!(\log, v_1, w_1)$ of $s_{\mathbf{b}}$, and $t_\tau^2 \stackrel{\text{def}}{=} \mathbf{a}?v_1.\mathbf{a}?v_2.\tau.\mathbf{a}!w_2.\mathbf{b}!(\log, v_2, w_2)$ of $s_{\mathbf{c}}$, one can easily deduce that no other adequate action disabling monitor can enforce φ_1 with fewer modifications than those imposed by m_{det} , namely, $mc(m_{\text{det}}, t_\tau^1) = mc(m_{\text{det}}, t_\tau^2) = 1$. Furthermore, consider the invalid traces $t_\tau^1\{c/a\}$ and $t_\tau^2\{c/a\}$ that are respectively produced by versions of $s_{\mathbf{b}}$ and $s_{\mathbf{c}}$ that interact on some port c instead of \mathbf{a} (for any port $c \neq \mathbf{a}$). Since m_{det} binds the port c to its data binder x and uses this information in its insertion branch, $(\bullet, (x?(-))).Y$, the *same* modification count is achieved for these traces, as well i.e., $mc(m_{\text{det}}, t_\tau^1\{c/a\}) = mc(m_{\text{det}}, t_\tau^2\{c/a\}) = 1$. \blacksquare

Example 6.4 describes the case where formula φ is optimally enforced by a *finite-state* and *finitely-branching* monitor, m_{det} . In the general case, this is not always possible.

Example 6.5. Consider formula φ_2 stating that an initial input on port \mathbf{a} followed by another input from some other port $x_2 \neq \mathbf{a}$ constitutes invalid system behaviour. Also consider monitor m_1 where $\text{enf}(m_1, \varphi_2)$.

$$\begin{aligned} \varphi_2 &\stackrel{\text{def}}{=} [(\mathbf{a}?(-))][((x_2)?(-), x_2 \neq \mathbf{a})]\text{ff} \\ m_1 &\stackrel{\text{def}}{=} (\mathbf{a}?(-)).\text{rec } Y.((\bullet, \mathbf{b}?v_{\text{def}}).Y + (\mathbf{a}?(-)).\text{id}) \end{aligned}$$

When enforcing a system that generates the run $t_\tau^3 \stackrel{\text{def}}{=} \mathbf{a}?v_1.\mathbf{b}?v_2.\mathbf{a}!w_1.u_\tau^3$, monitor m_1 modifies the trace only *once*. Although it disables the input $\mathbf{b}?v_2$, it subsequently unblocks the SuS by inserting $\mathbf{b}?v_{\text{def}}$ and so trace t_τ^3 is transformed into $\mathbf{a}?v_1.\tau.\mathbf{a}!w_1.u_\tau^3$. However, for a slightly modified version of t_τ^3 , e.g., $t_\tau^3\{c/b\}$, m_1 scores a modification count of $2 + |u_\tau^3|$. This is the case because, although it blocks the invalid input on port c , it fails to insert the default value that unblocks the SuS. A more expressive version of m_1 , such as

$$m_2 \stackrel{\text{def}}{=} (\mathbf{a}?(-)).\text{rec } Y.((\bullet, \mathbf{b}?v_{\text{def}}).Y + \underline{(\bullet, \mathbf{c}?v_{\text{def}}).Y} + (\mathbf{a}?(-)).\text{id}),$$

circumvents this problem by defining an extra insertion branch (underlined), but still fails to be optimal in the case of $t_\tau^3\{d/b\}$. In this case, there does not exist a way to finitely define a monitor that can insert a default value on every possible input port $x_2 \neq \mathbf{a}$. Hence, it means that the optimal monitor m_{opt} for φ_1 would be an *infinite branching* one, i.e., it requires a countably infinite summation that is not expressible in TRN,

$$m_{\text{opt}} \stackrel{\text{def}}{=} (\mathbf{a}?(-)).(\text{rec } Y. \sum_{\mathbf{b} \in \text{PORT and } \mathbf{a} \neq \mathbf{b}} (\bullet, \mathbf{b}?v_{\text{def}}).Y + (\mathbf{a}?(-)).\text{id})$$

or alternatively

$$(\mathbf{a}?(-)).(\text{rec } Y. \sum_{\mathbf{b} \in \text{PORT}} (\bullet, \mathbf{a} \neq \mathbf{b}, \mathbf{b}?v_{\text{def}}).Y + (\mathbf{a}?(-)).\text{id})$$

where the condition $a \neq b$ is evaluated at runtime. \blacksquare

Unlike Example 6.4, Example 6.5 presents a case where optimality can only be attained by a monitor that defines an *infinite number of branches*; this is problematic since monitors are required to be *finitely* described. As it is not always possible to find a finite monitor that enforces a formula using the least amount of transformation for every possible system, this indicates that Definition 6.3 is too strict. We thus mitigate this issue by weakening Definition 6.3 and redefine it in terms of the set of system states SYS_Π , i.e., the set of states that can only perform inputs using the ports specified in a finite $\Pi \subset \text{PORT}$. Although this weaker version does *not* guarantee that the monitor m optimally enforces φ on *all* possible systems, it, however, ensures optimal enforcement for all the systems that input values via the ports specified in Π .

Definition 6.6 (Weak Optimal Enforcement). A monitor m is *weakly optimal* when enforcing φ , denoted as $\text{oenf}(m, \varphi, \Pi)$, iff it *enforces* φ (Definition 4.9) and when for every state $s \in \text{SYS}_\Pi$, explicit trace t_τ and monitor n , if $\text{ec}(n) \subseteq \text{ec}(m)$, $\text{enf}(n, \varphi)$ and $s \xrightarrow{t_\tau}$ then $\text{mc}(m, t_\tau) \leq \text{mc}(n, t_\tau)$. \square

Example 6.7. Monitor m_1 from Example 6.5 ensures that φ_2 is optimally enforced on systems that interact on ports a and b , i.e., when $\Pi = \{a, b\}$, while monitor m_2 guarantees it when $\Pi = \{a, b, c\}$. \blacksquare

We can show that a synthesised monitor $(\llbracket \varphi, \Pi \rrbracket)$ obtained using the synthesis function Definition 5.3 from Section 5 is also guaranteed to be weakly optimal (as stated by Definition 6.6) when enforcing φ on a SuS s whose input ports are specified by Π , i.e., $s \in \text{SYS}_\Pi$. Since our synthesis produces only action disabling monitors, i.e., $\text{ec}(\llbracket \varphi, \Pi \rrbracket) = \{\text{DIS}\}$ for all φ and Π , we can limit ourselves to monitors pertaining to the set $\text{DISTRN} \stackrel{\text{def}}{=} \{n \mid \text{if } \text{ec}(n) \subseteq \{\text{DIS}\}\}$. The proof for Theorem 6.12 below relies on the following lemmas.

Lemma 6.8. For every $m \in \text{DISTRN}$ and explicit trace t_τ , there exists some N such that $\text{mc}(m, t_\tau) = N$. \square

Lemma 6.9. For every action α and monitor $m \in \text{DISTRN}$, if it is the case that $m \xrightarrow{\alpha \blacktriangleright \alpha} m'$, $\text{enf}(m, \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i)$ and $(p_j, c_j)(\alpha) = \sigma$ (for some $j \in I$) then $\text{enf}(m', \varphi_j \sigma)$. \square

Lemma 6.10. For every monitor $m \in \text{DISTRN}$, whenever $\text{enf}(m, \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i)$ and, for some m' , $m \xrightarrow{(a!v) \blacktriangleright \bullet} m'$ then $\text{enf}(m', \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i)$. \square

Lemma 6.11. For every monitor $m \in \text{DISTRN}$, whenever $\text{enf}(m, \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i)$ and, for some m' , $m \xrightarrow{\bullet \blacktriangleright (a?v)} m'$ then $\text{enf}(m', \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i)$. \square

Theorem 6.12 (Weak Optimal Enforcement). For every system $s \in \text{SYS}_\Pi$, explicit trace t_τ and monitor m , if $\text{ec}(m) \subseteq \text{ec}(\llbracket \varphi, \Pi \rrbracket)$, $\text{enf}(m, \varphi)$ and $s \xrightarrow{t_\tau}$ implies $\text{mc}(\llbracket \varphi, \Pi \rrbracket, t_\tau) \leq \text{mc}(m, t_\tau)$.

Proof. Since, from Lemma 6.8, we know that for every $m \in \text{DISTRN}$, there exists some N such that $\text{mc}(m, t_\tau) = N$, we can prove that if $\text{enf}(m, \varphi)$, $s \xrightarrow{t_\tau}$ and $\text{mc}(\llbracket \varphi, \Pi \rrbracket, t_\tau) = N$ then $N \leq \text{mc}(m, t_\tau)$. We proceed by rule induction on $\text{mc}(\llbracket \varphi, \Pi \rrbracket, t_\tau)$.

Case $\text{mc}(\llbracket \varphi, \Pi \rrbracket, t_\tau)$ when $t_\tau = \mu t'_\tau$ and $(\llbracket \varphi, \Pi \rrbracket)[\text{sys}(\mu t'_\tau)] \xrightarrow{\mu} m'_\varphi[\text{sys}(t'_\tau)]$. Assume that

$$\text{mc}(\llbracket \varphi, \Pi \rrbracket, \mu t'_\tau) = \text{mc}(m'_\varphi, t'_\tau) = N \quad (6.1)$$

which implies that

$$\langle \varphi, \Pi \rangle [\mathbf{sys}(\mu t'_\tau)] \xrightarrow{\mu} m'_\varphi [\mathbf{sys}(t'_\tau)] \quad (6.2)$$

and also assume that

$$\mathbf{enf}(m, \varphi) \quad (6.3)$$

and that $s \xrightarrow{\mu t'_\tau}$. By the rules in our model we can infer that the reduction in (6.2) can result from rule \mathbf{IASY} when $\mu = \tau$, \mathbf{IDeF} and \mathbf{ITRNO} when $\mu = \mathbf{a!}v$, or \mathbf{ITRNI} when $\mu = \mathbf{a?}v$. We consider each case individually.

- \mathbf{IASY} : By rule \mathbf{IASY} from (6.2) we know that $\mu = \tau$ and that

$$m'_\varphi = \langle \varphi, \Pi \rangle. \quad (6.4)$$

Since from (6.3) we know that m is sound and eventual transparent, we can thus deduce that m does not hinder internal τ -actions from occurring and so the composite system $\langle \varphi, \Pi \rangle [\mathbf{sys}(\tau t'_\tau)]$ can always transition over τ via rule \mathbf{IASY} , that is,

$$m[\mathbf{sys}(\tau t'_\tau)] \xrightarrow{\tau} m[\mathbf{sys}(t'_\tau)]. \quad (6.5)$$

Hence, by (6.1), (6.3) and since $s \xrightarrow{\tau t'_\tau}$ entails $s \xrightarrow{\tau} s'$ and $s' \xrightarrow{t'_\tau}$ we can apply the *inductive hypothesis* and deduce that $N \leq mc(m, t'_\tau)$ so that by (6.5) and the definition of mc , we conclude that $N \leq mc(m, \tau t'_\tau)$ as required.

- \mathbf{IDeF} : From (6.2) and rule \mathbf{IDeF} we know that $\mu = \mathbf{a!}v$, $\langle \varphi, \Pi \rangle \xrightarrow{\mathbf{a!}v}$ and that $m'_\varphi = \mathbf{id}$. Since \mathbf{id} does not modify actions, we can deduce that $mc(m'_\varphi, t'_\tau) = 0$ and so by the definition of mc we know that $mc(\langle \varphi, \Pi \rangle, (\mathbf{a!}v)t'_\tau) = 0$ as well. This means that we cannot find a monitor that performs fewer transformations, and so we conclude that $0 \leq mc(m, (\mathbf{a!}v)t'_\tau)$ as required.
- \mathbf{ITRNI} : From (6.2) and rule \mathbf{ITRNI} we know that $\mu = \mathbf{a?}v$ and that

$$\langle \varphi, \Pi \rangle \xrightarrow{(\mathbf{a?}v)\blacktriangleright(\mathbf{a?}v)} m'_\varphi. \quad (6.6)$$

We now inspect the cases for φ .

- $\varphi \in \{\mathbf{ff}, \mathbf{tt}, X\}$: The cases for \mathbf{ff} and X do not apply since $\langle \mathbf{ff}, \Pi \rangle$ and $\langle X, \Pi \rangle$ do not yield a valid monitor, while the case when $\varphi = \mathbf{tt}$ gets trivially satisfied since $\langle \mathbf{tt}, \Pi \rangle = \mathbf{id}$ and $mc(\mathbf{id}, (\mathbf{a?}v)t'_\tau) = 0$.
- $\varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$ where $\#_{i \in I}(p_i, c_i)$: Since $\varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$, by the definition of $\langle - \rangle$ we have that

$$\begin{aligned} & \langle \varphi \wedge = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i, \Pi \rangle \\ &= \mathbf{rec} Y. \left(\sum_{i \in I} \begin{cases} \mathbf{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \mathbf{ff} \\ (p_i, c_i). \langle \varphi_i, \Pi \rangle & \text{otherwise} \end{cases} \right) + \mathbf{def}(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i) \\ &= \left(\sum_{i \in I} \begin{cases} \mathbf{dis}(p_i, c_i, \langle \varphi \wedge =, \Pi \rangle, \Pi) & \text{if } \varphi_i = \mathbf{ff} \\ (p_i, c_i). \langle \varphi_i, \Pi \rangle & \text{otherwise} \end{cases} \right) + \mathbf{def}(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i) \end{aligned} \quad (6.7)$$

Since normalized conjunctions are disjoint, i.e., $\#_{i \in I}(p_i, c_i)$, from (6.7) we can infer that the identity reduction in (6.6) can only happen when $\mathbf{a?}v$ matches an identity branch,

$(p_j, c_j).(\varphi_j, \Pi)$ (for some $j \in I$), and so we have that

$$(p_j, c_j)(a?v) = \sigma. \quad (6.8)$$

Hence, knowing (6.6) and (6.8), by rule ETRN we know that $m'_\varphi = (\varphi_j\sigma, \Pi)$ and so by (6.1) we can infer that

$$mc(m'_\varphi, t'_\tau) = N \text{ where } m'_\varphi = (\varphi_j\sigma, \Pi). \quad (6.9)$$

Since from (6.7) we also know that the monitor branch $(p_j, c_j).(\varphi_j, \Pi)$ is derived from a non-violating modal necessity, i.e., $[(p_j, c_j)]\varphi_j$ where $\varphi_j \neq \text{ff}$, we can infer that $a?v$ is not a violating action and so it should not be modified by any other monitor m , as otherwise it would infringe the eventual transparency constraint of assumption (6.3). Therefore, we can deduce that

$$m \xrightarrow{(a?v)\blacktriangleright(a?v)} m' \quad (\text{for some } m') \quad (6.10)$$

and subsequently, knowing (6.10) and that $t_\tau = (a?v)t'_\tau$ and also that $\text{sys}((a?v)t'_\tau) \xrightarrow{a?v} \text{sys}(t'_\tau)$, by rule ITRNI and the definition of mc we infer that

$$mc(m, (a?v)t'_\tau) = mc(m', t'_\tau). \quad (6.11)$$

As by (6.3), (6.6), (6.8) and Lemma 6.9 we know that $\text{enf}(m', \varphi_j\sigma)$, by (6.9) and since $s \xrightarrow{(a?v)t'_\tau}$ entails that $s \xrightarrow{a?v} s'$ and $s' \xrightarrow{t'_\tau}$, we can apply the *inductive hypothesis* and deduce that $N \leq mc(m', t'_\tau)$ and so from (6.11) we conclude that $N \leq mc(m, (a?v)t'_\tau)$ as required.

- $\varphi = \max X.\varphi'$ and $X \in \mathbf{fv}(\varphi')$: Since $\varphi = \max X.\varphi'$, by the syntactic restrictions of $\text{SHML}_{\mathbf{nf}}$ we infer that φ' cannot be ff or tt since $X \notin \mathbf{fv}(\varphi')$ otherwise, and it cannot be X since every logical variable must be guarded. Hence, φ' must be of a specific form, i.e., $\max Y_1 \dots Y_n. \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$, and so by unfolding every fixpoint in $\max X.\varphi'$ we reduce our formula to $\varphi \stackrel{\text{def}}{=} \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i \{ \max X.\varphi' / X, \dots \}$. We thus omit the remainder of this proof as it becomes identical to that of the subcase when $\varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$.

- ITRNO: We elide the proof for this case as it is very similar to that of ITRNI.

Case $mc((\varphi, \Pi), t_\tau)$ when $t_\tau = \mu t'_\tau$ and $(\varphi, \Pi)[\text{sys}(\mu t'_\tau)] \xrightarrow{\mu'} m'_\varphi[\text{sys}(t'_\tau)]$ and $\mu' \neq \mu$.
Assume that

$$mc((\varphi, \Pi), \mu t'_\tau) = 1 + M \quad (6.12)$$

$$\text{where } M = mc(m'_\varphi, t'_\tau) \quad (6.13)$$

which implies that

$$(\varphi, \Pi)[\text{sys}(\mu t'_\tau)] \xrightarrow{\mu'} m'_\varphi[\text{sys}(t'_\tau)] \text{ where } \mu' \neq \mu \quad (6.14)$$

and also assume that

$$\text{enf}(m, \varphi) \quad (6.15)$$

and that $s \xrightarrow{\mu t'_\tau}$. Since we only consider action disabling monitors, the μ' reduction of (6.14) can only be achieved via rules `iDisO` or `iDisI`. We thus explore both cases.

- `iDisI`: From (6.14) and by rule `iDisI` we have that $\mu = a?v$ and $\mu' = \tau$ and that

$$\langle \varphi, \Pi \rangle \xrightarrow{\bullet \blacktriangleright a?v} m'_\varphi. \quad (6.16)$$

We now inspect the cases for φ .

- $\varphi \in \{\text{ff}, \text{tt}, X\}$: These cases do not apply since $\langle \text{ff}, \Pi \rangle$ and $\langle X, \Pi \rangle$ do not yield a valid monitor, while $\langle \text{tt}, \Pi \rangle = \text{id}$ does not perform the reduction in (6.16).
- $\varphi = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$ where $\#_{i \in I}(p_i, c_i)$: Since $\varphi = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i$, by the definition of $\langle - \rangle$ we have that

$$\begin{aligned} & \langle \varphi_\wedge = \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle \\ &= \text{rec } Y. \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i). \langle \varphi_i, \Pi \rangle & \text{otherwise} \end{cases} \right) + \text{def}(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i) \\ &= \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, \langle \varphi_\wedge, \Pi \rangle, \Pi) & \text{if } \varphi_i = \text{ff} \\ (p_i, c_i). \langle \varphi_i, \Pi \rangle & \text{otherwise} \end{cases} \right) + \text{def}(\bigwedge_{i \in I} [(p_i, c_i)] \varphi_i) \end{aligned} \quad (6.17)$$

Since normalized conjunctions are disjoint *i.e.*, $\#_{i \in I}(p_i, c_i)$, and since $s \xrightarrow{\mu t'_\tau}$ where $\mu = (a?v)$, by the definition of `dis`, from (6.17) we deduce that the reduction in (6.16) can only be performed by an insertion branch of the form, $(\bullet, c_j \{a/x\}, a?v). \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle$ that can only be derived from a violating modal necessity $[(p_j, c_j)]\text{ff}$ (for some $j \in I$). Hence, we can infer that

$$m'_\varphi = \langle \bigwedge_{i \in I} [(p_i, c_i)] \varphi_i, \Pi \rangle \quad (6.18)$$

$$p_j = (x)?(y) \text{ and } c_j \{a/x\} \Downarrow \text{true}. \quad (6.19)$$

Knowing (6.19) and that $[(p_j, c_j)]\text{ff}$ we can deduce that any input on port a is erroneous and so for the soundness constraint of assumption (6.15) to hold, any other monitor m is obliged to somehow *block* this input port. As we consider action disabling monitors, *i.e.*, $m \in \text{DisTrn}$, we can infer that monitor m may block this input in two ways, namely, either by not reacting to the input action, *i.e.*, $m \not\xrightarrow{a?v}$, or by additionally inserting a default value v , *i.e.*, $m \xrightarrow{\bullet \blacktriangleright (a?v)} m'$. We explore both cases.

* $m \not\xrightarrow{a?v}$: Since $\text{sys}((a?v)t'_\tau) \xrightarrow{a?v} \text{sys}(t'_\tau)$ and since $m \not\xrightarrow{a?v}$, by the rules in our model we know that for every action μ' , $m[\text{sys}((a?v)t'_\tau)] \not\xrightarrow{\mu'}$ and so by the definition of `mc` we have that $mc(m, (a?v)t'_\tau) = |(a?v)t'_\tau|$ meaning that by blocking inputs on a , m also blocks (and thus modifies) every subsequent action of trace t'_τ . Hence, this suffices to deduce that *at worst* $1 + M$ is equal to $|(a?v)t'_\tau|$, that is $1 + M \leq |(a?v)t'_\tau|$, and so from (6.12) we can deduce that $1 + M \leq mc(\langle \varphi, \Pi \rangle, \mu t'_\tau)$ as required.

* $m \xrightarrow{\bullet \blacktriangleright (a?v)} m'$: Since $\text{sys}((a?v)t'_\tau) \xrightarrow{a?v} \text{sys}(t'_\tau)$ and since $m \xrightarrow{\bullet \blacktriangleright (a?v)} m'$, by rule `iDisI` we know that $m[\text{sys}((a?v)t'_\tau)] \xrightarrow{\tau} m[\text{sys}(t'_\tau)]$ and so by the definition of `mc` we have that

$$mc(m, (a?v)t'_\tau) = 1 + mc(m', t'_\tau). \quad (6.20)$$

As by (6.15), (6.16) and Lemma 6.11 we infer that $\text{enf}(m', \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i)$, by (6.13), (6.20) and since $s \xrightarrow{(a?v)t'_\tau}$ entails that $s \xrightarrow{(a?v)} s'$ and $s' \xrightarrow{t'_\tau}$, we can apply the *inductive hypothesis* and deduce that $M \leq mc(m', t'_\tau)$ and so from (6.12), (6.13) and (6.20) we conclude that $1 + M \leq mc(m, (a?v)t'_\tau)$ as required.

- $\varphi = \max X.\varphi'$ and $X \in \mathbf{fv}(\varphi')$: We omit showing this proof as it is a special case of when $\varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$.
- **iDISO**: We omit showing the proof for this subcase as it is very similar to that of case **iDISI**. Apart from the obvious differences (e.g., $a!v$ instead of $a?v$), Lemma 6.10 is used instead of Lemma 6.11.

Case $mc(\langle \varphi, \Pi \rangle, t_\tau)$ when $t_\tau \in \{\mu t'_\tau, \varepsilon\}$ and $\langle \varphi, \Pi \rangle[\text{sys}(\mu t'_\tau)] \not\xrightarrow{\mu'}$. Assume that

$$mc(\langle \varphi, \Pi \rangle, t_\tau) = |t_\tau| \quad (\text{where } t_\tau \in \{\mu t'_\tau, \varepsilon\}) \quad (6.21)$$

$$\langle \varphi, \Pi \rangle[\text{sys}(\mu t'_\tau)] \not\xrightarrow{\mu'} \quad (6.22)$$

$$\text{enf}(m, \varphi) \quad (6.23)$$

Since $t_\tau \in \{\mu t'_\tau, \varepsilon\}$ we consider both cases individually.

- $t_\tau = \varepsilon$: This case holds trivially since by (6.21), (6.22) and the definition of mc , $mc(\langle \varphi, \Pi \rangle, \varepsilon) = |\varepsilon| = 0$.
- $t_\tau = \mu t'_\tau$: Since $t_\tau = \mu t'_\tau$ we can immediately exclude the cases when $\mu \in \{\tau, a!v\}$ since rules **iASY** and **iDEF** make it impossible for (6.22) to be attained in such cases. Particularly, rule **iASY** always permits the SuS to independently perform an internal τ -move, while rule **iDEF** allows the monitor to default to **id** whenever the system performs an unspecified output $a!v$. However, in the case of inputs, $a?v$, the monitor may completely block inputs on a port a and as a consequence cause the entire composite system $\langle \varphi, \Pi \rangle[\text{sys}(\mu t'_\tau)]$ to block, thereby making (6.22) a possible scenario. We thus inspect the cases for φ vis-a-vis $\mu = a?v$.
 - $\varphi \in \{\mathbf{ff}, \mathbf{tt}, X\}$: These cases do not apply since $\langle \mathbf{ff}, \Pi \rangle$ and $\langle X, \Pi \rangle$ do not yield a valid monitor and since $\langle \mathbf{tt}, \Pi \rangle = \mathbf{id}$ is incapable of attaining (6.22).
 - $\varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$ where $\#_{i \in I}(p_i, c_i)$: Since $\varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$, by the definition of $\langle - \rangle$ we have that

$$\begin{aligned} & \langle \varphi \wedge, \Pi \rangle \\ &= \text{rec } Y. \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, Y, \Pi) & \text{if } \varphi_i = \mathbf{ff} \\ (p_i, c_i). \langle \varphi_i, \Pi \rangle & \text{otherwise} \end{cases} \right) + \text{def} \left(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i \right) \\ &= \left(\sum_{i \in I} \begin{cases} \text{dis}(p_i, c_i, \langle \varphi \wedge, \Pi \rangle, \Pi) & \text{if } \varphi_i = \mathbf{ff} \\ (p_i, c_i). \langle \varphi_i, \Pi \rangle & \text{otherwise} \end{cases} \right) + \text{def} \left(\bigwedge_{i \in I} [(p_i, c_i)]\varphi_i \right) \end{aligned} \quad (6.24)$$

Since $\mu = a?v$, from (6.24) and by the definitions of **dis** and **def** we can infer that the only case when (6.22) is possible is when the inputs on port a satisfy a violating modal necessity, that is, there exists some $j \in I$ such that $[(p_j, c_j)]\mathbf{ff}$ and for every $v \in \text{VAL}$, $\text{match}(p_j, a?v) = \sigma$ and $c_j \sigma \Downarrow \mathbf{true}$. At the same time, the monitor is also *unaware* of the port on which the erroneous input can be made, i.e., $a \notin \Pi$. Hence, this case does not apply since we limit ourselves to SYS_Π , i.e., states of system that can only input values via the ports specified in Π .

- $\varphi = \max X.\varphi'$: As argued in previous cases, this is a special case of $\varphi = \bigwedge_{i \in I} [(p_i, c_i)]\varphi_i$ and so we omit this part of the proof.

Case $mc(\llbracket \varphi, \Pi \rrbracket, t_\tau)$ when $t_\tau \in \{\mu t'_\tau, \varepsilon\}$ and $\llbracket \varphi, \Pi \rrbracket[\mathbf{sys}(t_\tau)] \xrightarrow{\mu'} m'_\varphi[\mathbf{sys}(t_\tau)]$. As we only consider action disabling monitors, this case does not apply since the transition $\llbracket \varphi, \Pi \rrbracket[\mathbf{sys}(t_\tau)] \xrightarrow{\mu'} m'_\varphi[\mathbf{sys}(t_\tau)]$ can only be achieved via action enabling and rules IENO and IENI . \square

7. CONCLUSIONS AND RELATED WORK

This work extends the framework presented in the precursor to this work [ACFI18] to the setting of bidirectional enforcement where observable actions such as inputs and outputs require different treatment. We achieve this by:

- (1) augmenting substantially our instrumentation relation (Figure 4);
- (2) refining our definition of enforcement to incorporate transparency over violating systems (Definition 4.9);
- (3) providing a more extensive synthesis function (Definition 5.3) that is proven correct (Theorem 5.5); and
- (4) exploring notions of transducer optimality in terms of limited levels of intrusiveness (Definitions 6.3 and 6.6 and Theorem 6.12).

Future work. There are a number of possible avenues for extending our work. One immediate step would be the implementation of the monitor operational model presented in Section 3 together with the synthesis function described in Section 5. This effort should be integrated it within the detectEr tool suite [CFS15, AF16, CFAI17, CFA⁺17, AAA⁺21]. This would allow us to assess the overhead induced by our proposed bidirectional monitoring [AAFI21]. Another possible direction would be the development of behavioural theories for the transducer operational model presented in Section 3, along the lines of the refinement preorders studied in earlier work on sequence recognisers [Fra21, Fra17, AAF⁺21a]. Finally, applications of this theory on transducers, along the lines of [LMM20] are also worth exploring.

Related work. In his seminal work [Sch00], Schneider introduced the concept of runtime enforcement and enforceability in a linear-time setting. Particularly, in his setting a property is deemed enforceable if its *violation* can be *detected* by a *truncation automaton*, and prevented via system termination. By preventing misbehaviour, these automata can only enforce safety properties. Ligatti *et al.* extended this work in [LBW05] via *edit automata*—an enforcement mechanism capable of *suppressing* and *inserting* system actions. A property is thus enforceable if it can be expressed as an edit automaton that *transforms* invalid executions into valid ones via suppressions and insertions. As a means to assess the correctness of these automata, the authors introduced *soundness* and *transparency*.

Both settings by Schneider [Sch00] and Ligatti *et al.* [LBW05] assume a trace based view of the SuS and that every action can be freely manipulated by the monitor. They also do not distinguish between the specification and the enforcement mechanism, as properties are encoded in terms of the enforcement model itself, *i.e.*, as edit/truncation automata. In our prior work [ACFI18], we addressed this issue by separating the specification and verification aspects of the logic and explored the enforceability of μHML in a unidirectional

context and in relation to a definition of adequate enforcement defined in terms of soundness and transparency. In this paper we adopt a stricter notion of enforceability that requires adherence to eventual transparency and investigate the enforceability of sHML formulas in the context of bidirectional enforcement.

Bielova and Massacci [Bie11, BM11b] remark that, on their own, soundness and transparency fail to specify the extent in which a transducer should modify invalid runtime behaviour and thus introduce a *predictability* criterion. A transducer is *predictable* if one can predict the edit-distance between an invalid execution and a valid one. With this criterion, adequate monitors are further restricted by setting an upper bound on the number of transformations that a monitor can apply to correct invalid traces. Although this is similar to our notion of optimality, we however use it to compare an adequate (sound and eventual transparent) monitor to *all* the other adequate monitors and determine whether it is the least intrusive monitor that can enforce the property of interest.

In [KAB⁺17] Könighofer *et al.* present a synthesis algorithm similar to our own that produces action replacement monitors called *shields* from safety properties encoded as automata-based specifications. Although their shields can analyse both the inputs and outputs of a reactive system, they still perform unidirectional enforcement since they only modify the data associated with the system’s output actions. By definition, shields should adhere to correctness and minimum deviation which are, in some sense, analogous to soundness and transparency respectively.

In [PRS⁺16, PRS⁺17], Pinisetty *et al.* conduct a preliminary investigation of RE in a bidirectional setting. They, however, model the behaviour of the SuS as a trace of input and output pairs, *a.k.a. reactions*, and focus on enforcing properties by modifying the payloads exchanged by these reactions. This way of modelling system behaviour is, however, quite restrictive as it only applies to synchronous reactive systems that output a value in reaction to an input. This differs substantially from the way we model systems as LTSs, particularly since we can model more complex systems that may opt to collect data from multiple inputs, or supply multiple outputs in response to an input. The enforcement abilities studied in [PRS⁺16, PRS⁺17] are also confined to action replacement that only allows the monitor to modify the data exchanged by the system in its reactions, and so the monitors in [PRS⁺16, PRS⁺17] are unable to disable and enable actions. Due to their trace based view of the system, their correctness specifications do not allow for defining correct system behaviour in view of its different execution branches. This is particularly useful when considering systems whose inputs may lead them into taking erroneous computation branches that produce invalid outputs. Moreover, since their systems do not model communication ports, their monitors cannot influence directly the control structure of the SuS, *e.g.*, by opening, closing or rerouting data through different ports.

Finally, Lanotte *et al.* [LMM20] employ similar synthesis techniques and correctness criteria to ours (Definitions 4.2 and 4.4) to generate enforcement monitors for a timed setting. They apply their process-based approach to build tools that enforce data-oriented security properties. Although their implementations handle the enforcement of first-order properties, the theory on which it is based does not, nor does it investigate logic enforceability.

REFERENCES

- [AAA⁺21] Duncan Paul Attard, Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir, and Karoliina Lehtinen. Better late than never or: Verifying asynchronous components at runtime. In Kirstin Peters and Tim A. C. Willemse, editors, *Formal Techniques for Distributed Objects*,

- Components, and Systems - 41st IFIP WG 6.1 International Conference, FORTE 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14-18, 2021, Proceedings*, volume 12719 of *Lecture Notes in Computer Science*, pages 207–225. Springer, 2021.
- [AAA⁺22] Luca Aceto, Antonis Achilleos, Duncan Paul Attard, Léo Exibard, Adrian Francalanza, and Anna Ingólfssdóttir. A monitoring tool for linear-time μ hml. In Maurice H. ter Beek and Marjan Sirjani, editors, *Coordination Models and Languages - 24th IFIP WG 6.1 International Conference, COORDINATION 2022, Held as Part of the 17th International Federated Conference on Distributed Computing Techniques, DisCoTec 2022, Lucca, Italy, June 13-17, 2022, Proceedings*, volume 13271 of *Lecture Notes in Computer Science*, pages 200–219. Springer, 2022.
- [AAF⁺19] Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir, and Karoliina Lehtinen. Adventures in monitorability: From branching to linear time and back again. *Proc. ACM Program. Lang.*, 3(POPL):52:1–52:29, January 2019.
- [AAF⁺20] Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir, and Sævar Örn Kjartansson. Determinizing monitors for HML with recursion. *J. Log. Algebraic Methods Program.*, 111:100515, 2020.
- [AAF⁺21a] Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir, and Karoliina Lehtinen. The Best a Monitor Can Do. In *CSL*, volume 183 of *LIPICs*, pages 7:1–7:23. Schloss Dagstuhl, 2021.
- [AAF⁺21b] Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir, and Karoliina Lehtinen. An operational guide to monitorability with applications to regular properties. *Softw. Syst. Model.*, 20(2):335–361, 2021.
- [AAFI18a] Luca Aceto, Antonis Achilleos, Adrian Francalanza, and Anna Ingólfssdóttir. A Framework for Parameterized Monitorability. In *FoSSaCS*, pages 203–220. Springer, 2018.
- [AAFI18b] Luca Aceto, Antonis Achilleos, Adrian Francalanza, and Anna Ingólfssdóttir. Monitoring for silent actions. In Satya Lokam and R. Ramanujam, editors, *FSTTCS 2017: Foundations of Software Technology and Theoretical Computer Science*, volume 93 of *LIPICs*, pages 7:1–7:14. Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [AAFI21] Luca Aceto, Duncan Paul Attard, Adrian Francalanza, and Anna Ingólfssdóttir. On benchmarking for concurrent runtime verification. In *FASE*, volume 12649 of *LNCS*, pages 3–23. Springer, 2021.
- [ACFI18] Luca Aceto, Ian Cassar, Adrian Francalanza, and Anna Ingólfssdóttir. On Runtime Enforcement via Suppressions. In *CONCUR*, volume 118, pages 34:1–34:17. Schloss Dagstuhl, 2018.
- [ACFI21] Luca Aceto, Ian Cassar, Adrian Francalanza, and Anna Ingólfssdóttir. On bidirectional runtime enforcement. In Kirstin Peters and Tim A. C. Willemse, editors, *Formal Techniques for Distributed Objects, Components, and Systems - 41st IFIP WG 6.1 International Conference, FORTE 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14-18, 2021, Proceedings*, volume 12719 of *Lecture Notes in Computer Science*, pages 3–21. Springer, 2021.
- [AEF⁺22] Antonis Achilleos, Léo Exibard, Adrian Francalanza, Karoliina Lehtinen, and Jasmine Xuereb. A synthesis tool for optimal monitors in a branching-time setting. In Maurice H. ter Beek and Marjan Sirjani, editors, *Coordination Models and Languages - 24th IFIP WG 6.1 International Conference, COORDINATION 2022, Held as Part of the 17th International Federated Conference on Distributed Computing Techniques, DisCoTec 2022, Lucca, Italy, June 13-17, 2022, Proceedings*, volume 13271 of *Lecture Notes in Computer Science*, pages 181–199. Springer, 2022.
- [AF16] Duncan Paul Attard and Adrian Francalanza. A Monitoring Tool for a Branching-Time Logic. In *RV*, pages 473–481. Springer, 2016.
- [AI99] Luca Aceto and Anna Ingólfssdóttir. Testing Hennessy-Milner Logic with Recursion. In Wolfgang Thomas, editor, *FoSSaCS*, pages 41–55. Springer, 1999.
- [AILS07] Luca Aceto, Anna Ingólfssdóttir, Kim Guldstrand Larsen, and Jiri Srba. *Reactive Systems: Modelling, Specification and Verification*. Cambridge Univ. Press, NY, USA, 2007.
- [Av11] Rajeev Alur and Pavol Černý. Streaming Transducers for Algorithmic Verification of Single-pass List-processing Programs. In *POPL*, pages 599–610. ACM, 2011.
- [BAF21] Matthew Alan Le Brun, Duncan Paul Attard, and Adrian Francalanza. Graft: general purpose raft consensus in elixir. In Stavros Aronis and Annette Bieniusa, editors, *Proceedings of the 20th*

- ACM SIGPLAN International Workshop on Erlang, Erlang@ICFP 2021, Virtual Event, Korea, August 26, 2021*, pages 2–14. ACM, 2021.
- [BCD⁺17] Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. *TCS*, 669:33 – 58, 2017.
- [Bie11] Nataliia Bielova. *A theory of constructive and predictable runtime enforcement mechanisms*. PhD thesis, University of Trento, 2011.
- [BM11a] Nataliia Bielova and Fabio Massacci. Do you really mean what you actually enforced?-edited automata revisited. *Journal of Information Security*, 10(4):239–254, 2011.
- [BM11b] Nataliia Bielova and Fabio Massacci. Predictability of enforcement. In *International Symposium on Engineering Secure Software and Systems*, pages 73–86. Springer, 2011.
- [CBD⁺12] Tzu-Chun Chen, Laura Bocchi, Pierre-Malo Deniérou, Kohei Honda, and Nobuko Yoshida. Asynchronous Distributed Monitoring for Multiparty Session Enforcement. In *TGC*, pages 25–45. Springer, 2012.
- [CFA⁺17] Ian Cassar, Adrian Francalanza, Duncan Paul Attard, Luca Aceto, and Anna Ingólfssdóttir. A Suite of Monitoring Tools for Erlang. In *RV-CuBES*, volume 3 of *Kalpa Publications in Computing*, pages 41–47. EasyChair, 2017.
- [CFAI17] Ian Cassar, Adrian Francalanza, Luca Aceto, and Anna Ingólfssdóttir. eAOP: An Aspect Oriented Programming Framework for Erlang. In *Erlang*, ACM SIGPLAN, 2017.
- [CFS15] Ian Cassar, Adrian Francalanza, and Simon Said. Improving Runtime Overheads for detectEr. In *FESCA*, volume 178 of *EPTCS*, pages 1–8, 2015.
- [FAI17] Adrian Francalanza, Luca Aceto, and Anna Ingólfssdóttir. Monitorability for the Hennessy-Milner logic with recursion. *Formal Methods in System Design*, 51(1):87–116, 2017.
- [FFM08] Yliès Falcone, Jean-Claude Fernandez, and Laurent Mounier. Synthesizing enforcement monitors wrt. the safety-progress classification of properties. In *Information Systems Security*. Springer, 2008.
- [FFM12] Yliès Falcone, Jean-Claude Fernandez, and Laurent Mounier. What can you verify and enforce at runtime? *Journal on Software Tools for Technology Transfer*, 14(3):349, 2012.
- [Fra17] Adrian Francalanza. Consistently-Detecting Monitors. In *CONCUR*, volume 85 of *LIPICs*, pages 8:1–8:19. Dagstuhl, Germany, 2017.
- [Fra21] Adrian Francalanza. A Theory of Monitors. *Information and Computation*, page 104704, 2021.
- [HL95] M. Hennessy and X. Liu. A Modal Logic for Message Passing Processes. *Acta Informatica*, 32(4):375–393, Apr 1995.
- [HL96] M. Hennessy and H. Lin. Proof systems for message-passing process algebras. *Formal Aspects of Computing*, 8(4):379–407, 1996.
- [HRTZ18] Klaus Havelund, Giles Reger, Daniel Thoma, and Eugen Zălinescu. *Monitoring Events that Carry Data*, pages 61–102. Springer International Publishing, Cham, 2018.
- [JGP16] Limin Jia, Hannah Gommerstadt, and Frank Pfenning. Monitors and blame assignment for higher-order session types. In *POPL*, pages 582–594, NY, USA, 2016. ACM.
- [KAB⁺17] Bettina Könighofer, Mohammed Alshiekh, Roderick Bloem, Laura Humphrey, Robert Könighofer, Ufuk Topcu, and Chao Wang. Shield synthesis. *Formal Methods in System Design*, 51(2):332–361, Nov 2017.
- [Koz83] Dexter C. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [KT12] Raphaël Khoury and Nadia Tawbi. Which security policies are enforceable by runtime monitors? a survey. *Computer Science Review*, 6(1):27–45, 2012.
- [Lar90] Kim G Larsen. Proof systems for satisfiability in Hennessy-Milner logic with recursion. *Theoretical Computer Science*, 72(2):265–288, 1990.
- [LBW05] Jay Ligatti, Lujo Bauer, and David Walker. Edit automata: enforcement mechanisms for run-time security policies. *Journal of Information Security*, 4(1):2–16, 2005.
- [LBW09] Jay Ligatti, Lujo Bauer, and David Walker. Run-time enforcement of nonsafety policies. *ACM Trans. Inf. Syst. Secur.*, 12(3):19:1–19:41, January 2009.
- [LMM20] Ruggero Lanotte, Massimo Merro, and Andrei Munteanu. Runtime enforcement for control system security. In *CSF*, pages 246–261. IEEE, 2020.
- [MPW92] Robin Milner, Joachim Parrow, and David Walker. A Calculus of Mobile Processes, I. *Information and Computation*, 100(1):1–40, 1992.

- [PRS⁺16] Srinivas Pinisetty, Partha S. Roop, Steven Smyth, Stavros Tripakis, and Reinhard von Hanxleden. Runtime enforcement of reactive systems using synchronous enforcers. *CoRR*, abs/1612.05030, 2016.
- [PRS⁺17] Srinivas Pinisetty, Partha S. Roop, Steven Smyth, Nathan Allen, Stavros Tripakis, and Reinhard Von Hanxleden. Runtime enforcement of cyber-physical systems. *ACM Trans. Embed. Comput. Syst.*, 16(5s):178:1–178:25, September 2017.
- [RH97] J. Rathke and M. Hennessy. Local model checking for value-passing processes (extended abstract). In *TACS*, pages 250–266. Springer, 1997.
- [Sak09] Jacques Sakarovitch. *Elements of Automata Theory*. Cambridge University Press, New York, NY, USA, 2009.
- [San11] Davide Sangiorgi. *Introduction to Bisimulation and Coinduction*. Cambridge University Press, New York, NY, USA, 2011.
- [Sch00] Fred B Schneider. Enforceable security policies. *ACM Transactions on Information and System Security (TISSEC)*, 3(1):30–50, 2000.
- [vHRF17] A. C. van Hulst, M. A. Reniers, and W. J. Fokkink. Maximally permissive controlled system synthesis for non-determinism and modal logic. *Discrete Event Dynamic Systems*, 27(1):109–142, 2017.