

Consistently-Detecting Monitors*

Adrian Francalanza

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

Abstract

We study a contextual definition for deterministic monitoring based on consistent detections. It is defined in terms of the observed behaviour of the monitor when instrumented over arbitrary systems. We give an alternative, coinductive definition based on controllability which does not rely on system quantifications, and show that it is fully-abstract wrt. the former definition. We then develop a symbolic counterpart to the controllability definition to facilitate an automated analysis for controllable monitors involving data.

1998 ACM Subject Classification F.3.1 Specifying and Verifying and Reasoning about Programs, F.3.2 Semantics of Programming Languages - Process models, D.2.5 Testing and Debugging

Keywords and phrases Runtime Monitoring, Deterministic Behaviour, Controllability, Compositional Reasoning, Symbolic Analysis

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2017.8

1 Introduction

Monitors are computational entities that observe the executions of other entities (referred to hereafter as *systems*) with the aim of accruing system information [32, 26], comparing system executions against some behavioural specification [23, 7], or reacting to the observed executions via adaptation or enforcement procedures [13, 36]. They are typically considered to be part of the *Trusted Computing Base (TCB)* and, consequently, their descriptions are expected to be *correct*. A correctness requirement often presumed of monitors is that they should exhibit *deterministic behaviour*. Yet, for most monitoring frameworks, such a requirement is seldom specified in unambiguous terms. In fact, there are a number of viable alternatives that one could consider (e.g., [44, 29, 25, 1]) and it is unclear how to choose one over the other in an objective manner. Moreover, these definitions often fail to account for the instrumentation mechanism used to compose a monitor with the system under scrutiny which may, in turn, affect monitoring behaviour. All of this leads to a poor understanding of what should be expected of a monitor, and may give rise to discrepancies between these expectations and what needs to be guaranteed by the monitor implementer in practice.

Non-determinism is intrinsic to a number of computational models used for expressing monitors and monitored systems. In fact, a substantial body of work on monitors is either cast in terms of inherently non-deterministic formalisms such as Büchi automata [45, 17], or formalisms that admit non-determinism such as process calculi and labelled transition systems [12, 46, 30, 22, 10, 23]. Non-deterministic computation arises naturally in concurrent and distributed programming, used increasingly for runtime monitoring [37, 24, 21, 8, 11]. Furthermore, a growing number of monitoring tools employ automata-based specification languages [16, 5, 41, 18] that offer rudimentary support for ensuring deterministic behaviour:

* Partly supported by UoM and RANNIS projects CPSRP05-04 and TheoFoMon:163406-051.



their respective implementations are either thread-unsafe [41] or admit arbitrary code for transition-triggered actions [14, 16].

This paper sets out to investigate deterministic behaviour for monitors. The study is limited to execution monitors (sequence recognisers) [43, 36], used extensively for Runtime Verification (RV). Our work is developed in terms of an expository formalism (similar to the aforementioned work on transition-based descriptions) expressing monitored systems that can analyse trace events carrying data and admit degrees of non-determinism. We propose a *contextual* definition for deterministic monitor behaviour, founded on the *observational* behaviour that can be discerned when a monitor is instrumented to execute with any arbitrary system under scrutiny. The definition serves two purposes. First, its contextual nature allows us to admit *as many* correct monitors *as possible*, as long as these cannot be *externally* perceived to behave non-deterministically—we contend that the resulting definition is fairly intuitive. Second, it allows us to *justify* design decisions for an alternative definition describing the deterministic behaviour of monitors, based instead on the notion of *controllability* [31]. We show a correspondence between these two definitions. In addition, we demonstrate how the alternative definition (which is arguably less intuitive than its contextual counterpart) is more amenable to automated analyses for assessing the deterministic behaviour exhibited by monitors. In particular, we study how this alternative definition can be reformulated in symbolic terms, to facilitate a tractable handling of infinite-state monitor analysis due to data.

► **Example 1.** The monitor description m_1 accepts traces from an authenticator, that challenges (event `chl`) a supplicant for an *arbitrary* value x followed by the supplicant’s authentication (`aut`) with the (correct) encoding of x , $y = \text{enc}(x)$. The authenticator subsequently acknowledges (`ack`) using the same value y . The guard construct $\text{chl}(x)$ quantifies over *any* value x ; guards $\text{aut}\langle v \rangle$ and $\text{ack}\langle v \rangle$ require authentications (*resp.* acknowledgments) for a *specific* value v . \top denotes acceptance.

$$\begin{aligned} m_1 &\triangleq \text{chl}(x).\text{let } y = \text{enc}(x) \text{ in } (\text{aut}\langle y \rangle.\text{ack}\langle y \rangle.\top) \\ m_2 &\triangleq \text{chl}(x).\text{let } y = \text{enc}(x) \text{ in } (\text{aut}\langle y \rangle.\text{ack}\langle y \rangle.\top + \text{aut}\langle z \rangle.\text{if } z \neq y \text{ then } \text{ack}\langle z' \rangle.\perp) \\ m_3 &\triangleq \text{chl}(x).\text{let } y = \text{enc}(x) \text{ in } (\text{aut}\langle y \rangle.\text{ack}\langle y \rangle.\top + \text{aut}\langle z \rangle.\text{if } z \neq y \text{ then } \text{ack}\langle z' \rangle.\perp \text{ else } \text{ack}\langle z \rangle.\top) \end{aligned}$$

It is easy to inadvertently introduce non-determinism. Monitor m_2 extends m_1 (using the choice operator, $+$) with the intention of rejecting (note the verdict \perp) acknowledgments following authentications whose value is *not* the encoding of the challenge value x . When $v_2 \neq \text{enc}(v_1)$, the violating trace $t = \text{chl}\langle v_1 \rangle.\text{aut}\langle v_2 \rangle.\text{ack}\langle v_2 \rangle \dots$ is always rejected by m_2 . More subtly, however, when $v_2 = \text{enc}(v_1)$ the trace t may cause the monitor to non-deterministically choose either branch, whereby the unintended branch does *not* reach a \top verdict. But non-determinism is not necessarily conducive to inconsistent verdicts and, in cases where verdicts are considered as the only monitor observable behaviour, such non-determinism may be tolerated. Monitor m_3 *deterministically accepts* trace t when $v_2 = \text{enc}(v_1)$, albeit along different execution paths, and *deterministically rejects* it whenever $v_2 \neq \text{enc}(v_1)$. ◀

The main contributions of the paper are:

1. A contextual definition for observationally deterministic monitoring behaviour, Definition 6.
2. An alternative definition based on controllability, Definition 11, that coincides with it, Theorem 13.
3. Symbolic variants of the latter definition enabling tractable automations, Theorems 21 and 25.

Monitors

	$w, o \in \text{VERD} ::= \top$	(accept)	$ \perp$	(reject)
	$ \mathbf{0}$	(inconclusive)		
	$m, n \in \text{MON} ::= w$	(verdict)	$ \text{let } x = e \text{ in } m$	(evaluate)
	$ l\langle e \rangle.m$	$(\text{expression guard})$	$ l(x).m$	$(\text{quantified guard})$
	$ m + n$	(choice)	$ \text{if } b \text{ then } m \text{ else } n$	(conditional)
	$ \text{rec } X.m$	(recursion)	$ X$	$(\text{monitor variable})$
VER	IFT	IFF	LET	
$\frac{}{w \xrightarrow{\eta} w}$	$\frac{\llbracket b \rrbracket = \text{true}}{\text{if } b \text{ then } m \text{ else } n \xrightarrow{\tau} m}$	$\frac{\llbracket b \rrbracket = \text{false}}{\text{if } b \text{ then } m \text{ else } n \xrightarrow{\tau} n}$	$\frac{\llbracket e \rrbracket = v}{\text{let } x = e \text{ in } m \xrightarrow{\tau} m[v/x]}$	
GRE	GRQ	REC	CH1	
$\frac{\llbracket e \rrbracket = v}{l\langle e \rangle.m \xrightarrow{l\langle v \rangle} m}$	$\frac{}{l(x).m \xrightarrow{l\langle v \rangle} m[v/x]}$	$\frac{}{\text{rec } X.m \xrightarrow{\tau} m[\text{rec } X.m/X]}$	$\frac{m \xrightarrow{\alpha} m'}{m + n \xrightarrow{\alpha} m'}$	

Instrumentation

MON	TER	ASs	ASm
$\frac{s \xrightarrow{\eta} r \quad m \xrightarrow{\eta} n}{s \triangleleft m \xrightarrow{\eta} r \triangleleft n}$	$\frac{s \xrightarrow{\eta} r \quad m \not\xrightarrow{\eta} \quad m \not\xrightarrow{\tau}}{s \triangleleft m \xrightarrow{\eta} r \triangleleft \mathbf{0}}$	$\frac{s \xrightarrow{\tau} r}{s \triangleleft m \xrightarrow{\tau} r \triangleleft m}$	$\frac{m \xrightarrow{\tau} n}{s \triangleleft m \xrightarrow{\tau} s \triangleleft n}$

■ **Figure 1** A Model for describing Instrumented Systems.

The paper is structured as follows. Section 2 presents the monitor framework used for this study, allowing us motivate our touchstone definition for deterministic monitor behaviour in Section 3. Section 4 presents a fully-abstract alternative definition that is amenable to compositional reasoning. In Sections 5 and 6 symbolic variants are developed for automation purposes. Section 7 concludes.

2 Systems, Monitors, Instrumentation and Monitored Systems

We perceive *systems* as entities that generate *events* while executing. *Observable* events, $\eta \in \text{EVT}$, are those visible to a monitor and have the form $l\langle v \rangle$ where l is an event *label* taken from a set $l, k \in \text{LAB}$, and v is an event *payload* taken from some unspecified value domain $v, u \in \text{VAL}$. Events capture a number of computational notions such as inputs/outputs in message-passing programs [22], or method/function calls and returns [13, 38]. To simplify our technical development, we consider monadic (*i.e.*, single-valued) events but the formalism can be extended to accommodate polyadicity.

Systems may be described as Labelled Transition Systems (LTSs) [2, 34], triples $\langle \text{SYS}, \text{ACT}, \longrightarrow \rangle$ consisting of a set of states, $s, r \in \text{SYS}$, a set of actions, $\alpha, \beta \in \text{ACT} = \text{EVT} \cup \{\tau\}$ that include all observable events EVT and a distinguished silent action $\tau \notin \text{EVT}$ for *unobservable* events, and a transition relation, $\longrightarrow \subseteq (\text{SYS} \times \text{ACT} \times \text{SYS})$. The suggestive notation $s \xrightarrow{\alpha} s'$ denotes $(s, \alpha, s') \in \longrightarrow$, whereas $s \not\xrightarrow{\alpha} s'$ denotes $\neg(\exists s' \cdot s \xrightarrow{\alpha} s')$. As usual, we write $s \Longrightarrow s'$ in lieu of $s(\xrightarrow{\tau})^* s'$ and $s \xRightarrow{\eta} s'$ for $s \Longrightarrow \cdot \xrightarrow{\eta} \cdot \Longrightarrow s'$, referring to s' as a η -*derivative* of s ; $s \xRightarrow{t} s'$ stands for $\exists s' \cdot s \xRightarrow{t} s'$. We let *traces*, $t, u \in \text{EVT}^*$, range over (finite) sequences of *observable* events and write $s \xRightarrow{\eta_1} \dots \xRightarrow{\eta_n} s_n$ as $s \xRightarrow{t} s_n$, where $t = \eta_1, \dots, \eta_n$. The notation $u \dots$ is occasionally used to denote the existence of some trace t with a prefix u .

We presuppose an expression language $e, d \in \text{EXP}$ that ranges over the (event) value domain VAL and a denumerable set of *expression variables* $x, y, z \in \text{VARS}$; \vec{e} and \vec{x} resp.

denote lists of expressions and variables. We also assume a boolean expression language $b, c \in \text{BEXP}$ defined over EXP that includes standard constructs for conjunctions, $b \wedge c$, and negations, $\neg b$, but also equality predicates over expressions $e = d$. The meta-functions $\mathbf{fv}(e)$ and $\mathbf{fv}(b)$ return the *free variables* in the resp. expressions; expressions are *closed* whenever $\mathbf{fv}(e) = \emptyset$ and *open* otherwise, and similarly for boolean expressions. *Valuations* are total maps from variables to values, $\rho \in (\text{VARS} \rightarrow \text{VAL})$ whereas *substitutions* are partial maps from variables to expressions $\sigma \in (\text{VARS} \rightarrow \text{EXP})$; substitutions are denoted as $[\vec{e}/\vec{x}]$, where $d[\vec{e}/\vec{x}]$ represents the (simultaneous) substitution of all occurrences of $x_i \in \vec{x}$ in d by the corresponding $e_i \in \vec{e}$. We assume an *evaluation function* that takes an expression and a valuation and returns a value, $\llbracket e \rho \rrbracket = v$. Similarly, boolean expressions have a semantic function mapping them to the boolean domain via a valuation, $\llbracket b \rho \rrbracket \in \{\text{true}, \text{false}\}$, where we presume the expected properties, e.g., $\llbracket (b \wedge c) \rho \rrbracket = \text{true}$ iff $\llbracket b \rho \rrbracket = \text{true}$ and $\llbracket c \rho \rrbracket = \text{true}$. We also assume a classical interpretation of boolean expressions, i.e., $\llbracket (\neg b) \rho \rrbracket = \text{true}$ iff $\llbracket b \rho \rrbracket = \text{false}$. To alleviate the presentation, we often work up to associativity and commutativity for conjunctions, treating $(\text{BEXP}, \wedge, \text{true})$ as an abelian monoid. For closed expressions, we elide the valuation and write $\llbracket e \rrbracket$ and $\llbracket b \rrbracket$ for $\llbracket e \rho \rrbracket$ and $\llbracket b \rho \rrbracket$ resp. The *satisfiability* judgement for boolean expressions, $\text{sat}(b) \stackrel{\text{def}}{=} \exists \rho. \llbracket b \rho \rrbracket = \text{true}$, plays a central role in subsequent development.

Monitors, here defined by the syntax in Figure 1, constitute the focus of our study. They may reach two kinds of verdicts, VERD . *Conclusive verdicts* consist of acceptances, \top , and rejections, \perp . In addition, a monitor may also reach the *inconclusive verdict*, $\mathbf{0}$, a form of premature termination used when the generated system events of the monitor specification itself does not yield sufficient information so as to reach a definite conclusion. The monitor *expression* guard $l\langle e \rangle.m$ expects events with label l and a payload value matching the evaluation of e , whereas the *quantified* guard $l(x).m$ allows the monitor to *dynamically learn* the payload of an event with label l . Monitors may branch (externally) depending on the events observed, $m+n$, or branch (internally) based on data predicates, *if b then m else n*. They may also perform internal computation themselves by evaluating expressions, *let x = e in m*, or recurse, $\text{rec } X.p$, via *term variables* $X, Y, Z \in \text{TVARS}$. The constructs $l(x).m$ and $\text{rec } X.m$ act as *binders* for x and X resp. in m , inducing the usual notions of open/closed (monitor) terms. We work up to alpha-conversion of bound expression/term variables and use the shorthand *if b then m* for *if b then m else 0* and $\tau.m$ for *let x = v in m* where $x \notin \mathbf{fv}(m)$.

The semantics of *closed* monitors is also defined in terms of an LTS, via the transition rules in Figure 1: for each $m \in \text{MON}$ we have a dedicated LTS $\langle M, \text{ACT}, \longrightarrow \rangle$ where $M \subseteq \text{MON}$ are the monitors reachable from m via transitions. The rules model the monitor analysis of observable events. Rule VER describes how verdicts are *irrevocable*, meaning that a verdict can analyse *any* observable event but always transition to itself. In rule GRE , an expression guard $l\langle e \rangle.m$ only transitions to the continuation m when observing an event matching the label l with the payload equal to $\llbracket e \rrbracket$. By contrast, a quantified guard $l(x).m$ transitions by analysing any event with label l , binding x to the event payload v in the continuation, $m[v/x]$; see rule GRQ . The remaining rules are as expected where the term $m[\text{rec } X.m/X]$ denotes the term substitution of $\text{rec } X.m$ for free occurrences of X in m .

A system s *instrumented* with a monitor m is referred to as a *monitored system* and denoted as $s \triangleleft m$. The semantics of monitored systems is defined by the instrumentation rules in Figure 1. We here adopt the composition relation studied in [22, 23], even though other instrumentation relations could have been used. Note that the chosen composition relation is still quite general: it is parametric wrt. the system and monitor abstract LTSs and it is *largely independent* of their specific language specifications, since it only requires the monitor LTS to contain an inconclusive (persistent) verdict state, $\mathbf{0}$. The instrumentation

relation of Figure 1 is *asymmetric*: a monitored system can transition with an observable event *only when* the system can produce that event *i.e.*, monitors are *passive* and cannot instigate transitions. When the system generates an (observable) event that can be analysed by the monitor, the two transition in lockstep according to their respective LTSs (rule MON). When the monitor *cannot* analyse the event generated ¹ and cannot internally transition to a state that enables it to do so (*i.e.*, it is already stable, $m \not\stackrel{\tau}{\rightarrow}$), the instrumentation does *not* block the monitored system: instead, it allows the system to transition *but* aborts monitoring to the inconclusive verdict (rule TER). System-monitor synchronisations are limited to observable events, and the specific entities can transition independently *wrt.* their respective internal moves (rules ASS and ASM).

► **Example 2.** Monitor m_4 below listens for input and output events $\text{in}\langle v \rangle$ and $\text{out}\langle v \rangle$ where the (integer) payload $v \in \mathbb{N}$ reports the port number over which the communication operation is performed.

$$m_4 \triangleq \text{rec } X. \left((\text{out}\langle 80 \rangle. \perp) + (\text{in}\langle x \rangle. \text{if } x=80 \text{ then } \text{out}\langle 81 \rangle. \top \text{ else } \text{out}\langle x \rangle. X) \right) \quad (1)$$

The monitor rejects system executions starting with an output on port 80 but accepts traces containing an input on port 80 followed by an output on port 81, preceded by an arbitrary number of input-output operations on any matching port other than 80. The execution below shows an *accepted* monitored computation for a system s generating the trace $\text{in}\langle 85 \rangle \cdot \text{out}\langle 85 \rangle \cdot \text{in}\langle 80 \rangle \cdot \text{out}\langle 81 \rangle$. In monitor m_4 , the binding on $\text{in}\langle x \rangle$ acts as a *freeze-variable* [19] for the subsequent $\text{out}\langle x \rangle$ guard in the else branch.

$$\begin{aligned} s \triangleleft m_4 &\stackrel{\tau}{\rightarrow} s \triangleleft (\text{out}\langle 80 \rangle. \perp) + (\text{in}\langle x \rangle. \text{if } x=80 \text{ then } \text{out}\langle 81 \rangle. \top \text{ else } \text{out}\langle x \rangle. m_4) && \text{REC} \\ &\xrightarrow{\text{in}\langle 85 \rangle} s' \triangleleft \text{if } 85=80 \text{ then } \text{out}\langle 81 \rangle. \top \text{ else } \text{out}\langle 85 \rangle. m_4 && \text{ChR} + \text{GrQ} \\ &\stackrel{\tau}{\rightarrow} s' \triangleleft \text{out}\langle 85 \rangle. m_4 \xrightarrow{\text{out}\langle 85 \rangle} s'' \triangleleft m_4 \xrightarrow{\text{in}\langle 80 \rangle \cdot \text{out}\langle 81 \rangle} s''' \triangleleft \top && \text{IFF}, \dots \end{aligned}$$

The instrumentation of Figure 1 delays system transitions to allow the monitor to internally transition to a state that can process the event. *E.g.*, if a system r can generate event $\text{out}\langle 80 \rangle$, $r \triangleleft m_4$ postpones this transition (MON and TER cannot be applied) until m_4 unfolds.

$$r \triangleleft m_4 \stackrel{\tau}{\rightarrow} r \triangleleft (\text{out}\langle 80 \rangle. \perp) + (\text{in}\langle x \rangle. \text{if } x=80 \text{ then } \dots) \xrightarrow{\text{out}\langle 80 \rangle} r' \triangleleft \perp \quad \text{ASM, MON}$$

Rule TER is crucial both for allowing monitored computations to proceed when the monitor cannot analyse an event, but also to avoid *unintended detections*. *E.g.*, if system r can generate the trace $\text{out}\langle 90 \rangle \cdot \text{in}\langle 80 \rangle \cdot \text{out}\langle 81 \rangle$, this behaviour should still be permitted when instrumented with the monitor m_4 , *but* the behaviour should not be detected according to the description in Equation 1. After the initial unfolding of m_4 , TER allows r to transition with $\text{out}\langle 90 \rangle$ but transitions m_4 to the inconclusive state, $\mathbf{0}$, since neither guard $\text{out}\langle 80 \rangle$ nor guard $\text{in}\langle x \rangle$ can process the event.

$$\begin{aligned} r \triangleleft m_4 &\stackrel{\tau}{\rightarrow} r \triangleleft (\text{out}\langle 80 \rangle. \perp) + (\text{in}\langle x \rangle. \text{if } x=80 \text{ then } \dots) \\ &\xrightarrow{\text{out}\langle 90 \rangle} r'' \triangleleft \mathbf{0} \xrightarrow{\text{in}\langle 80 \rangle \cdot \text{out}\langle 81 \rangle} r''' \triangleleft \mathbf{0} && \text{ASM, TER}, \dots \end{aligned}$$

Had rule TER been designed otherwise (leaving the monitor state unaltered when transiting with $\text{out}\langle 90 \rangle$) the ensuing events $\text{in}\langle 80 \rangle \cdot \text{out}\langle 81 \rangle$ would lead to the unintended acceptance of the trace. ◀

¹ This may be due to a number of reasons, such as event knowledge gaps or knowledge disagreements [6].

3 Deterministic Monitoring Behaviour

In a monitored system, non-deterministic behaviour can be caused by either the system or the monitor. We focus on identifying *non-determinism attributed to monitors*, teasing it apart from non-determinism caused by system behaviour. This is motivated by the fact that, generally, one has limited control over the behaviour of a system under scrutiny. We target a definition that admits monitor non-determinism that is not externally observable. Concretely, we consider detections (*i.e.*, conclusive verdicts) as the only externally visible aspect of a monitor and base our definition on the notion of *deterministic detections*—in applications such as RV, detections are associated with property satisfactions and violations [23, 7]. This immediately rules out a number of candidate definitions for deterministic monitor behaviour. For instance, a definition that considers a monitor m to be deterministic whenever, for *all* systems s and traces t , $s \triangleleft m \xrightarrow{t} s' \triangleleft m'$ and $s \triangleleft m \xrightarrow{t} s' \triangleleft m''$ implies $m' = m''$ is too stringent: it precludes the monitor description below (a slight modification on m_1 of Example 1)

$$m_5 \triangleq \text{chl}\langle x \rangle. \left((\text{let } y = \text{enc}(x) \text{ in aut}\langle y \rangle.\text{ack}\langle y \rangle.\top) + (\text{aut}\langle \text{enc}(x) \rangle.\text{ack}\langle \text{enc}(x) \rangle.\top) \right) \quad (2)$$

even though m_5 deterministically accepts traces of the form $\text{chl}\langle v_1 \rangle \cdot \text{aut}\langle v_2 \rangle \cdot \text{ack}\langle v_2 \rangle$ where $v_2 = \text{enc}(v_1)$. In fact, after an event $\text{chl}\langle v \rangle$ (for some value v), monitor m_5 can reach two possible internal states, namely $(\text{let } y = \text{enc}(v) \text{ in aut}\langle y \rangle.\text{ack}\langle y \rangle.\top) + (\text{aut}\langle \text{enc}(v) \rangle.\text{ack}\langle \text{enc}(v) \rangle.\top)$ or $\text{aut}\langle v' \rangle.\text{ack}\langle v' \rangle.\top$ where $v' = \text{enc}(v)$. Other candidates (*e.g.*, confluence defined over transitions [25, 40]) are either inadequate or not immediately applicable because they do not account for executions that do *not* lead to detections. *E.g.*, m_6 (below) would *not* be confluent (consider event $\text{in}\langle 81 \rangle$), even though it consistently rejects any trace with the prefix $u = \text{in}\langle 80 \rangle$ (and consistently does *not* detect all the other traces).

$$m_6 \triangleq (\text{in}\langle 80 \rangle.\perp) + ((\text{in}\langle 81 \rangle.\text{out}\langle 81 \rangle.\mathbf{0}) + (\text{in}\langle 81 \rangle.\text{out}\langle 81 \rangle.\text{in}\langle 82 \rangle.\mathbf{0})) \quad (3)$$

► **Definition 3** (Detected Computations). The transition sequence

$$s \triangleleft m \xrightarrow{t} s_0 \triangleleft m_0 \xrightarrow{\tau} s_1 \triangleleft m_1 \xrightarrow{\tau} s_2 \triangleleft m_2 \xrightarrow{\tau} \dots$$

is called a **t -computation** if it is **maximal** *i.e.*, either it is *infinite* or it is finite and *cannot be extended further* using τ -transitions. The t -computation above is called *accepted* whenever $\exists i \in \mathbb{N} \cdot m_i = \top$ and *rejected* when $\exists i \in \mathbb{N} \cdot m_i = \perp$. A *detected t -computation* is either an accepted or a rejected one. ◀

Detected computations are indexed by their trace to allow us to *partition* computations according to the system behaviour exhibited at runtime, thus accounting for *system non-determinism*. Definition 3 also permits monitors to stabilise and reach verdicts in the trailing τ -sequence following a t -trace.

► **Definition 4** (Deterministic Detection and Withholding). Monitor m *deterministically accepts* (*resp.* *deterministically rejects*) for system s along trace $t \in \text{EVT}^*$, denoted as $\text{da}(m, s, t)$ and $\text{dr}(m, s, t)$ *resp.*, iff *all* t -computation from $s \triangleleft m$ are accepting (*resp.* rejecting). Monitor m *deterministically detects* for s along t , $\text{dd}(m, s, t)$, whenever $\text{da}(m, s, t)$ or $\text{dr}(m, s, t)$. Monitor m *deterministically withholds* for s along trace t , $\text{dw}(m, s, t)$, iff *no* t -computation from $s \triangleleft m$ is accepting or rejecting. ◀

► **Example 5.** For arbitrary system s , monitors m_1 of Example 1 and m_5 of Equation 2 deterministically accept traces with the prefix $t = \text{chl}\langle v_1 \rangle.\text{aut}\langle v_2 \rangle.\text{ack}\langle v_2 \rangle$ where $v_2 = \text{enc}(v_1)$ and deterministically withhold on all the other traces. Monitor m_2 deterministically rejects traces with the prefix t above when $v_2 \neq \text{enc}(v_1)$ but does *not* deterministically detect traces with prefix t when $v_2 = \text{enc}(v_1)$. For arbitrary s , monitor m_3 deterministically detects *any* trace with the above prefix t (accepting or rejecting the trace depending on whether $v_2 = \text{enc}(v_1)$ or not) and deterministically withholds otherwise. For any system s , monitor m_4 of Example 2 satisfies $\text{dr}(m_4, s, t)$ when the trace t is of the form $t = \text{out}\langle 80 \rangle \dots, \text{da}(m_4, s, t)$ when $t = (\text{in}\langle v_i \rangle \cdot \text{out}\langle v_i \rangle)^i \cdot \text{in}\langle 80 \rangle \cdot \text{out}\langle 81 \rangle \dots$ for some $i \in \mathbb{N}$, and $\text{dw}(m_4, s, t)$ otherwise. Similarly, for all systems s , m_6 satisfies $\text{dr}(m, s, t)$ when $t = \text{in}\langle 80 \rangle \dots$ and $\text{dw}(m, s, t)$ otherwise. ◀

For the rest of our study, monitors with deterministic behaviour are defined as *consistently-detecting*.

► **Definition 6** (Consistent Detection). Monitor m *consistently detects* for system s , denoted as $\text{cd}(m, s)$ iff for *all* traces t we have $\text{dd}(m, s, t)$ or $\text{dw}(m, s, t)$. A monitor m is *consistently-detecting*, denoted as $\text{cd}(m)$, whenever $\text{cd}(m, s)$ holds for *any* system s . ◀

► **Example 7.** Monitors m_1, m_3, m_4, m_5 and m_6 are consistently-detecting, but m_2 is *not*. Definition 6 does *not* require monitors to perform any detections. The monitor $m_7 \triangleq \text{rec } X. (\text{in}\langle 81 \rangle.X) + (\text{in}\langle 81 \rangle.\text{out}\langle 81 \rangle.X)$ can consistently analyse an infinite number of traces for any s , $\text{cd}(m_7)$, even though it never flags. ◀

A few comments are in order. First, Definition 6 abstracts away from the particular instances of the systems considered, the specifics of the monitor language and instrumentation mechanism used; this makes it applicable to *arbitrary* monitoring setups. Second, $\text{cd}(m, s)$ may be seen as requiring an *ambiguity of $d=1$* from automata theory [29, 4], for the observable behaviour specified in Definition 4. Our setting is however more general, allowing for *infinite* states and alphabets (actions). Moreover, $\text{cd}(m)$ quantifies over *all* possible system compositions. Third, since Definition 6 is defined over *monitored system* behaviour, it allows us to assess the *actual* monitor behaviour at runtime. Particularly, the system quantification in $\text{cd}(m)$ accounts for any (indirect) effects of a system on the execution of a monitor.

► **Example 8.** Whereas monitor $m_8 \triangleq \text{in}\langle 81 \rangle.\perp$ is (trivially) consistently-detecting in the framework of Figure 1, the monitor $m_9 \triangleq \text{in}(x).\text{if } x = 81 \text{ then } \perp \text{ else } \mathbf{0}$ is, perhaps surprisingly, *not*. Consider a (diverging) system s with behaviour $s \xrightarrow{\text{in}\langle 81 \rangle} s' \xrightarrow{\tau} s'$. Although $s \triangleleft m_9$ can reject the t -computation for $t = \text{in}\langle 81 \rangle$, another possible t -computation of $s \triangleleft m_9$ is

$$s \triangleleft m_9 \xrightarrow{\text{in}\langle 81 \rangle} (s' \triangleleft \text{if } 81 = 81 \text{ then } \perp \text{ else } \mathbf{0}) \xrightarrow{\tau} (s' \triangleleft \text{if } 81 = 81 \text{ then } \perp \text{ else } \mathbf{0}) \xrightarrow{\tau} \dots$$

which *never* reaches a verdict. Therefore, we have $\neg \text{cd}(m_9)$ according to Definition 6. ◀

Fourth, consistently-detecting monitors are *not* compositional, affecting the subsequent machinery.

► **Example 9.** Although m_8 (from Example 8) and monitor $m_{10} = \text{in}\langle 81 \rangle.\top$ are both consistently-detecting according to Definition 6, their composition, *i.e.*, $m_8 + m_{10}$, is clearly *not*. ◀

4 Controllability

In spite of its generality and intuitive nature, Definition 6 it is hard to automate directly as a correctness analysis. One major obstacle is the inherent universal quantification over systems and traces defining $\text{cd}(m)$. In this section, we set out to give an alternative definition for describing consistently-detecting monitors that does not suffer from these shortcomings. It is based on the notion of *controllability* [20, 31] which, in discrete event settings, roughly refers to the ability to steer a (passive) entity to designated terminal states via a series of admissible controls. In our case, the monitors will constitute the passive entities to be steered, whereas the monitored systems assume the controller's role: the admissible controls are effectively the observable events in a monitoring setup that cause the monitor to transition, whereas the terminal states of interest are the conclusive verdicts. The proposed definition thus *inverts* the focus from how a system is monitored to how a monitor can be driven.

Before giving the actual definition, we first need to lift the technical machinery of Figure 1 to *sets* of monitors, $M, N \subseteq \text{MON}$: this allows us to express the status whereby a monitor that can be in a *number of potential states* after being driven by a sequence of steering controls, which facilitates the analysis of non-compositional properties such as ours (see Example 9).

► **Definition 10.** A monitor-set M *potentially reaches* a verdict w , $\text{pr}(M, w)$, when $\exists m \in M \cdot m \Rightarrow w$, and *potentially analyses* an event η , $\text{pa}(M, \eta)$, when $\exists m \in M \cdot m \xrightarrow{\eta}$. Function $\text{aft}(M, \eta)$ is defined as:

$$\begin{aligned} \text{aft}(M, \eta) &\stackrel{\text{def}}{=} \bigcup_{m \in M} \text{aft}(m, \eta) \\ \text{aft}(m, \eta) &\stackrel{\text{def}}{=} \{n \mid m \Rightarrow \cdot \xrightarrow{\eta} n\} \cup \{\mathbf{0} \mid \exists n \cdot m \Rightarrow n \not\xrightarrow{\eta} \text{ and } n \not\xrightarrow{\eta}\} \end{aligned}$$

◀

Intuitively, $\text{aft}(M, \eta)$ computes the set of reachable states from every $m \in M$ when it is asked by the instrumentation of Figure 1 to analyse an event η . The two conditions defining $\text{aft}(m, \eta)$ correspond to the monitored system transitions dictated by the respective rules MON and TER in Figure 1.

► **Definition 11 (Controllability).** A relation $\mathcal{R} \subseteq \mathcal{P}(\text{MON})$ is *controllable* iff for all $M \in \mathcal{R}$:

1. $\text{pr}(M, w)$ and $w \in \{\top, \perp\}$ implies $M = \{w\}$;
2. $\text{pa}(M, \eta)$ implies $\text{aft}(M, \eta) \in \mathcal{R}$.

Controllability, denoted as the relation \mathcal{C} , is the *largest* controllable relation. A monitor m (resp. monitor-set M) is said to be controllable iff $\{m\} \in \mathcal{C}$ (resp. $M \in \mathcal{C}$). ◀

Controllability is *coinductive*: to show that a monitor m is controllable, i.e., $\{m\} \in \mathcal{C}$, it suffices to provide a witness controllable relation \mathcal{R} such that $\{m\} \in \mathcal{R}$. Condition (i) in Definition 11 requires that if some $m \in M$ can reach a conclusive verdict, then *every* $m' \in M$ must be able to do so *immediately*, without requiring any preceding τ -moves (hence $M = \{w\}$); this rules out the possibility of inconsistent detections and, at the same time, prohibits diverging systems from interfering with the reaching of such verdicts (see Example 8). Condition (ii) in Definition 11 intuitively requires that this condition is satisfied for *any* event η observed, by *all* the states that any $m \in M$ may transition to when analysing η .

► **Example 12.** We can show that m_6 (Equation 3) is controllable via the controllable relation \mathcal{R}_1 below:

$$\mathcal{R}_1 = \left\{ \begin{array}{l} \{m_6\}, \{\top\}, \{\mathbf{0}, \text{in}\langle 82 \rangle.\mathbf{0}\}, \{\mathbf{0}\}, \\ \{\text{out}\langle 81 \rangle.\mathbf{0}, \text{out}\langle 81 \rangle.\text{in}\langle 82 \rangle.\mathbf{0}\} \end{array} \right\}$$

$$\mathcal{R}_2 = \left\{ \begin{array}{l} \{m_7\}, \{m_7, \text{out}\langle 81 \rangle.m_7\}, \{\mathbf{0}, m_7\}, \\ \{\mathbf{0}, m_7, \text{out}\langle 81 \rangle.m_7\}, \{\mathbf{0}\} \end{array} \right\}$$

Note that $\{m_6\} \in \mathcal{R}_1$. We can also *finitely* determine that the recursive monitor m_7 (Example 7) is controllable via the relation \mathcal{R}_2 . The reader may want to check that \mathcal{R}_2 is controllable. For instance,

$\mathbf{aft}(\{m_7, \text{out}\langle 81 \rangle.m_7\}, \text{in}\langle 81 \rangle) = \{\mathbf{0}, m_7, \text{out}\langle 81 \rangle.m_7\}$,
 $\mathbf{aft}(\{m_7, \text{out}\langle 81 \rangle.m_7\}, \text{out}\langle 81 \rangle) = \{\mathbf{0}, m_7\}$ and, importantly, $\mathbf{aft}(\{\mathbf{0}, m_7, \text{out}\langle 81 \rangle.m_7\}, \text{in}\langle 81 \rangle) = \{\mathbf{0}, m_7, \text{out}\langle 81 \rangle.m_7\}$ itself. ◀

Controllability coincides with Definition 6: we can use Definition 11 as a *sound* and *complete* proof technique to determine whether a monitor m satisfies $\text{cd}(m)$, side-stepping universal quantifications over systems.

► **Theorem 13** (Consistent Detection Full Abstraction). $\text{cd}(m)$ iff $\{m\} \in \mathcal{C}$ ◀

► **Example 14.** As a result of Theorem 13, we can show that m_6 and m_7 are *consistently-detecting* via the controllable relations \mathcal{R}_1 and \mathcal{R}_2 of Example 12. We can also *indirectly* show that $\neg \text{cd}(m_8 + m_{10})$ from Example 9, by arguing that there *cannot* be a controllable relation \mathcal{R} with $\{m_8 + m_{10}\} \in \mathcal{R}$. For suppose that such an \mathcal{R} exists. By Definition 11(ii) the monitor-set $\mathbf{aft}(\{m_8 + m_{10}\}, \text{in}\langle 81 \rangle) = \{\top, \perp\}$ must also be in \mathcal{R} ; this, in turn, would necessarily mean that \mathcal{R} is *not* controllable since $\{\top, \perp\}$ violates Definition 11(i). ◀

5 Symbolic Controllability

Controllability, Definition 11, is still not adequate for a fully automated analysis of consistently-detecting monitors. Particularly, whenever the *resp.* event value domain is infinite, quantified guards induce an infinite number of transitions, e.g., $l(x).m$ generates a transition with the label $l\langle v \rangle$ for every $v \in \text{VAL}$ (see rule GRQ). As a result, condition Definition 11(ii) may require the monitor analysis to consider a potentially infinite number of monitor-set states whenever monitor descriptions use quantified guards.

To this end, we define a *symbolic* semantics over *both* open² and closed monitor terms using *symbolic* events, $\theta \in \text{SEVT}$. These are similar to the events of Section 2 except that they carry *variables* instead of values as payloads, $l\langle x \rangle$. Symbolic transitions, $m \xrightarrow[b]{\mu} n$ are defined by the rules in Figure 2, where $\mu \in \text{SEVT} \cup \{\tau\}$ ranges over both symbolic and τ events, and the boolean expression b records the condition under which the LTS action may take place. For instance, the term $\text{rec } X.m$ may unfold in all circumstances (*i.e.*, $b = \text{true}$ in rule sREC) whereas the term $\text{if } b \text{ then } m \text{ else } n$ can either τ -transition to m when b holds, or to n when the converse, $\neg b$, holds (rules sIFT and sIFF). The other key rules in Figure 2 are sGRE and sGRQ: the former transitions with a symbolic event $l\langle x \rangle$ under the condition $x=e$, whereas the latter transitions with a similar symbolic event under *any* circumstance. Figure 2 also defines rules for *weak symbolic transitions*, $m \xrightarrow[b]{\theta} n$, and *reductions*, $m \Rightarrow_b n$,

² Open wrt. expression variables $x, y, \dots \in \text{VARS}$ not term variables $X, Y \dots \in \text{TVARS}$.

8:10 Consistently-Detecting Monitors

where both relations *aggregate* boolean constraints via conjunctions. Note that weak symbolic transitions describe transition sequences where τ -transitions *must precede* the (final) symbolic event transition. The predicate $m \xrightarrow{\mu} n$ denotes $\exists b, n \cdot m \xrightarrow{b} n$ whereas $m \xrightarrow{\theta} n$ stands for $\exists n \cdot m \xrightarrow{\theta} n$.

A *constrained monitor-set* $\langle b, M \rangle$ is a tuple where every $m \in M$ may be open, and b is a condition constraining free variables in M . Every $\langle b, M \rangle$ abstractly represents a (potentially infinite) set of *closed* monitor-sets for every valuation ρ satisfying b ,

$$\{ m\rho \mid \llbracket b\rho \rrbracket = \text{true} \text{ and } m \in M \} \quad (4)$$

In this sense, the monitor-sets in Section 4 are special cases of constrained monitor-sets where $b = \text{true}$ and M is closed. Note that whenever $\neg \text{sat}(b)$, the constrained monitor-set $\langle b, M \rangle$ denotes the *empty* set of monitor-sets, \emptyset , which is trivially controllable by Definition 11. We lift functions such as that for free variables $\mathbf{fv}(-)$ to constrained monitor-sets in the obvious manner, e.g., $\mathbf{fv}(\langle b, M \rangle) \stackrel{\text{def}}{=} \mathbf{fv}(b) \cup \mathbf{fv}(M)$.

► **Example 15.** The constrained monitor-set $\langle x \geq 3, \{\text{if } x = 2 \text{ then } \top \text{ else } \perp\} \rangle$ abstractly describes all monitor-sets $\{\text{if } x = 2 \text{ then } \top \text{ else } \perp\} \rho$ where $\rho(x) \geq 3$. For any such ρ , *no* monitor of the form $(\text{if } x = 2 \text{ then } \top \text{ else } \perp) \rho$ can transition to a \top verdict according to the concrete semantics of Figure 1. Symbolically, this may be expressed as $\neg \text{sat}(x \geq 3 \wedge x = 2)$. ◀

We abstractly model controllability, Definition 11, in terms of constrained monitor-sets, the symbolic semantics of Figure 2 and the satisfiability judgement $\text{sat}(b)$ defined earlier in Section 2.

► **Definition 16.** A constrained monitor-set $\langle b, M \rangle$ *potentially reaches* a verdict w , denoted as $\mathbf{spr}(\langle b, M \rangle, w)$, whenever $\exists m \in M, c \cdot m \xrightarrow{c} w$ and $\text{sat}(b \wedge c)$. Moreover, $\langle b, M \rangle$ *potentially analyses* a symbolic event θ along c , denoted as $\mathbf{spa}(\langle b, M \rangle, \theta, c)$, whenever $\exists m \in M \cdot m \xrightarrow{c} n$ and $\text{sat}(b \wedge c)$. ◀

Defining the symbolic counterpart to $\mathbf{aft}(M, \eta)$ of Definition 10 is less straightforward. Intuitively, from all the valuations ρ satisfying (and represented by) b in $\langle b, M \rangle$, only a *subset* of them may satisfy the condition c in a *potentially-analyses* judgement $\mathbf{spa}(\langle b, M \rangle, \theta, c)$ from Definition 16. A correct modelling of Definition 11 therefore requires us to take this fact into account.

► **Example 17.** Consider the constrained monitor-set $\langle \text{true}, M \rangle$ where $M = \{m_{10}, m_{11}\}$ and

$$m_{10} \triangleq \text{if } x=2 \text{ then } k\langle 1 \rangle.\top \text{ else } k\langle 1 \rangle.\perp \quad m_{11} \triangleq \text{if } x \leq 1 \vee x \geq 3 \text{ then } k\langle 1 \rangle.\perp \text{ else } k\langle 1 \rangle.\top$$

It turns out that for any ρ satisfying true , $M\rho$ is controllable. Therefore, for the judgement $\mathbf{spa}(\langle \text{true}, M \rangle, k\langle y \rangle, (x=2 \wedge y=1))$ of Definition 16, which holds since $m_{10} \xrightarrow{x=2 \wedge y=1} k\langle y \rangle$ and $\text{sat}(\text{true} \wedge x=2 \wedge y=1)$, the reachable states to be considered by a corresponding symbolic analysis (modelling *Definition 11(ii)*) should *not* include the residual state \perp , even though it may be reached after the event $k\langle y \rangle$ with $y = 1$ (see rule sGRE). The reason for this is that the conditions required to symbolically reach this state, i.e., $(\neg(x=2) \wedge y=1)$ or $((x \leq 1 \vee x \geq 3) \wedge y=1)$, *cannot* be satisfied by *any* ρ that also satisfies the $\mathbf{spa}(-)$ condition $(x=2 \wedge y=1)$. Symbolically, this may be expressed as $\neg \text{sat}((x=2 \wedge y=1) \wedge ((\neg(x=2) \wedge y=1)))$ and $\neg \text{sat}((x=2 \wedge y=1) \wedge ((x \leq 1 \vee x \geq 3) \wedge y=1))$. ◀

Symbolic Transitions

$$\begin{array}{c}
\text{sVER} \\
\frac{}{w \xrightarrow[\text{true}]{\eta} w} \\
\\
\text{sIFT} \\
\frac{}{\text{if } b \text{ then } m \text{ else } n \xrightarrow[b]{\tau} m} \\
\\
\text{sIFF} \\
\frac{}{\text{if } b \text{ then } m \text{ else } n \xrightarrow[\neg b]{\tau} n} \\
\\
\text{sREC} \\
\frac{}{\text{rec } X.m \xrightarrow[\text{true}]{\tau} m[\text{rec } X.m/X]} \\
\\
\text{sCH1} \\
\frac{m \xrightarrow[b]{\mu} m'}{m + n \xrightarrow[b]{\mu} m'} \\
\\
\text{sCH2} \\
\frac{n \xrightarrow[b]{\mu} n'}{m + n \xrightarrow[b]{\mu} n'} \\
\\
\text{sGRE} \\
\frac{}{l\langle e \rangle.m \xrightarrow[e=x]{l\langle x \rangle} m} \\
\\
\text{sGRQ} \\
\frac{}{l(y).m \xrightarrow[\text{true}]{l\langle x \rangle} m[x/y]} \\
\\
\text{sLET} \\
\frac{}{\text{let } x = e \text{ in } m \xrightarrow[\text{true}]{\tau} m[e/x]}
\end{array}$$

Weak Symbolic Transitions and Reductions

$$\begin{array}{c}
\text{sWTR1} \\
\frac{m \xrightarrow[b]{\theta} m'}{m \xrightarrow[b]{\theta} m'} \\
\\
\text{sWTR2} \\
\frac{m \xrightarrow[b]{\tau} m' \quad m' \xrightarrow[c]{\theta} m''}{m \xrightarrow[b \wedge c]{\theta} m''} \\
\\
\text{SWRD1} \\
\frac{}{m \xRightarrow[\text{true}]{} m} \\
\\
\text{sWTR2} \\
\frac{m \xrightarrow[b]{\tau} m' \quad m' \xRightarrow[c]{} m''}{m \xRightarrow[b \wedge c]{} m''}
\end{array}$$

■ **Figure 2** A Symbolic Semantics for Monitors.

The complications elicited in Example 17 are even more intricate. For instance, for a particular judgement $\mathbf{spa}(\langle b, M \rangle, \theta, c)$, one could have some $m_1, m_2 \in M$ whereby $m_i \xrightarrow[c_i]{\theta} m'_i$ and $\mathbf{sat}(b \wedge c \wedge c_i)$ for $i \in \{1, 2\}$, but at the same time having c_1 and c_2 being incompatible with one another, i.e., $\neg \mathbf{sat}(c_1 \wedge c_2)$. In such cases, the respective residual states m'_1 and m'_2 should be analysed separately.

► **Definition 18.** The *relevant conditions* for a monitor-set M wrt. a symbolic event θ are:

$$\mathbf{rc}(M, \theta) \stackrel{\text{def}}{=} \{c \mid \exists m \in M \cdot (m \xrightarrow[c]{\theta} \text{ or } \exists n \cdot (m \xRightarrow{c} n \text{ and } n \not\xrightarrow{\tau} \text{ and } n \not\xrightarrow{\theta}))\}$$

The *satisfiability combinations* for a condition-set $\{c_1, \dots, c_n\}$ wrt. a condition b are:

$$\mathbf{sc}(b, \{c_1, \dots, c_n\}) \stackrel{\text{def}}{=} \{\{b, c'_1, \dots, c'_n\} \mid \forall i \in 1..n \cdot (c'_i = c_i \text{ or } c'_i = \neg c_i)\}$$

The *reachable constrained monitor-sets* from $\langle b, M \rangle$ after θ with condition c are:

$$\begin{aligned}
\mathbf{saft}(\langle b, M \rangle, \theta, c) &\stackrel{\text{def}}{=} \{\langle \wedge B, \mathbf{saft}(M, B, \theta) \rangle \mid B \in \mathbf{sc}(b \wedge c, \mathbf{rc}(M, \theta)) \text{ and } \mathbf{sat}(\wedge B)\} \\
\mathbf{saft}(M, B, \theta) &\stackrel{\text{def}}{=} \left\{ n \mid \begin{array}{l} \exists m \in M, c \cdot \mathbf{sat}((\wedge B) \wedge c) \text{ and} \\ (m \xrightarrow[c]{\theta} n \text{ or } (\exists n' \cdot m \xRightarrow{c} n' \not\xrightarrow{\tau} \text{ and } n' \not\xrightarrow{\theta} \text{ and } n = \mathbf{0})) \end{array} \right\}
\end{aligned}$$

◀

In Definition 18, the *relevant conditions* for M wrt. θ , denoted as $\mathbf{rc}(M, \theta)$, are all the symbolic conditions that need to be considered to assess the reachable states from M for the symbolic event θ — they are the symbolic counterpart to the transition sequences defining $\mathbf{aft}(m, \eta)$ in Definition 10. The *satisfiability combinations* of a condition-set B wrt. a condition b , denoted as $\mathbf{sc}(b, B)$, capture the maximal condition subsets in B that any valuation ρ satisfying condition b also satisfies. Every condition set B' returned by $\mathbf{sc}(b, B)$

contains b itself and one condition c' for every boolean condition $c \in B$ (either c itself or its negation); these combination sets *partition* all the valuations ρ satisfying b . *Symbolic reachability* for $\langle b, M \rangle$ after θ with condition c , $\mathbf{saft}(\langle b, M \rangle, \theta, c)$ in Definition 18, is defined wrt. all the satisfiability combinations B of $\mathbf{rc}(M, \theta)$ for the (fixed) condition $b \wedge c$. Although $\mathbf{sc}(b \wedge c, \mathbf{rc}(M, \theta))$ partitions all the ρ satisfying $b \wedge c$, some of these partitions are *empty*. Accordingly, $\mathbf{saft}(\langle b, M \rangle, \theta, c)$ only considers the non-empty partitions via the satisfiability condition $\mathbf{sat}(\wedge B)$, where $\wedge B$ returns the syntactic conjunction formula $c_1 \wedge \dots \wedge c_n$ for a boolean set $B = \{c_1, \dots, c_n\}$.

It is worth remarking that the symbolic LTS of Figure 2, is image-finite [42], and thus finitely branching when considering the τ -transition graph of a term m . By König's Infinity Lemma [33] the set of constraints $\{c \mid m \xrightarrow{\theta}_c \text{ or } \exists n \cdot (m \xrightarrow{c} n \text{ and } n \not\xrightarrow{\tau} \text{ and } n \not\xrightarrow{\theta})\}$ must be finite and, as a result, $\mathbf{rc}(M, \theta)$ is finite too for a *finite* monitor-set M . This ensures that $\mathbf{saft}(\langle b, M \rangle, \theta, c)$ is well-defined.

► **Definition 19** (Symbolic Controllability). The relation $\mathcal{S} \subseteq (\text{BEXP} \times \mathcal{P}(\text{MON}))$ is called a *symbolically-controllable* relation iff for all constrained monitor-sets $\langle b, M \rangle \in \mathcal{S}$:

1. $\mathbf{spr}(\langle b, M \rangle, w)$ and $w \in \{\top, \perp\}$ implies $M = \{w\}$;
2. $\mathbf{spa}(\langle b, M \rangle, l\langle x \rangle, c)$ where $\mathbf{frsh}(\mathbf{fv}(\langle b, M \rangle)) = x$ implies $\mathbf{saft}(\langle b, M \rangle, l\langle x \rangle, c) \subseteq \mathcal{S}$.

Symbolic Controllability, denoted as \mathcal{C}_{sym} , is the largest symbolically-controllable relation. A (closed) monitor m is symbolically-controllable iff $\langle \text{true}, \{m\} \rangle \in \mathcal{C}_{\text{sym}}$. ◀

The clause Definition 19(ii) assumes a function $\mathbf{frsh}(V)$ that (deterministically) returns the *next* fresh variable x that is *not* in the variable set V . When compared to Definition 11(ii), this allows us to just consider *one* (symbolic) event, $l\langle x \rangle$, for a *finite* set of constraints, as opposed to a potentially infinite set of events, *i.e.*, $l\langle v \rangle$ for every $v \in \text{VAL}$.

► **Example 20.** Recall m_{10} and m_{11} from Example 17. The monitor $m_{12} \triangleq l\langle x \rangle.m_{10} + l\langle x \rangle.m_{11}$ can be shown to be symbolically-controllable via the relation \mathcal{S}_1 defined below, where $b_1 = (x = 2 \wedge y = 1)$, $b_2 = (\neg(x = 2) \wedge y = 1)$, $b_3 = ((x \leq 1 \vee x \geq 3) \wedge y = 1)$ and $b_4 = (\neg(x \leq 1 \vee x \geq 3) \wedge y = 1)$; these are obtained from the relevant conditions $\mathbf{rc}(\{m_{10}, m_{11}\}, k\langle y \rangle) = \{b_1, b_2, b_3, b_4\}$.

$$\mathcal{S}_1 = \left\{ \begin{array}{l} \langle \text{true}, \{m_{12}\} \rangle, \quad \langle \text{true}, \{m_{10}, m_{11}\} \rangle, \\ \langle (\text{true} \wedge b_1) \wedge b_1 \wedge \neg b_2 \wedge \neg b_3 \wedge b_4, \{\top\} \rangle, \quad \langle (\text{true} \wedge b_4) \wedge b_1 \wedge \neg b_2 \wedge \neg b_3 \wedge b_4, \{\top\} \rangle, \\ \langle (\text{true} \wedge b_2) \wedge \neg b_1 \wedge b_2 \wedge b_3 \wedge \neg b_4, \{\perp\} \rangle, \quad \langle (\text{true} \wedge b_3) \wedge \neg b_1 \wedge b_2 \wedge b_3 \wedge \neg b_4, \{\perp\} \rangle \end{array} \right\}$$

For illustrative purposes, we do not simplify the constraints in the constrained monitor-sets of \mathcal{S} to show how these are derived. *E.g.*, $\langle (\text{true} \wedge b_1) \wedge b_1 \wedge \neg b_2 \wedge \neg b_3 \wedge b_4, \{\top\} \rangle$ is obtained as a result of $\mathbf{saft}(\langle \text{true}, \{m_{10}, m_{11}\} \rangle, l\langle x \rangle, \text{true} \wedge b_1)$. In fact, the combination $\{(\text{true} \wedge b_1), b_1, \neg b_2, \neg b_3, b_4\}$ is the only satisfiable condition-set and all the others are filtered out by $\mathbf{saft}(\langle \text{true}, \{m_{10}, m_{11}\} \rangle, l\langle x \rangle, \text{true} \wedge b_1)$. ◀

Symbolic Controllability, Definition 19, is sound and complete wrt. Controllability, Definition 11.

► **Theorem 21** (Controllability Full Abstraction). $\{m\} \in \mathcal{C}$ iff $\langle \text{true}, \{m\} \rangle \in \mathcal{C}_{\text{sym}}$ ◀

► **Example 22.** Recall $m_3 \triangleq \text{chl}(x).m'_3$ from Example 1, recast in terms of m'_3 defined below as:

$$m'_3 \triangleq \text{let } y = \text{enc}(x) \text{ in } (\text{aut}\langle y \rangle.\text{ack}\langle y \rangle.\top + \text{aut}\langle z \rangle.\text{if } z \neq y \text{ then } \text{ack}\langle z' \rangle.\perp \text{ else } \text{ack}\langle z \rangle.\top)$$

Example 7 stated that m_3 is consistently-detecting. This fact is hard to determine using Definition 6, whereas analyses using Definition 11 are complicated by quantifications over

the values of events. By Theorems 13 and 21, we can show that m_3 is consistently-detecting via the *symbolic* controllability relation:

$$\mathcal{S}_2 = \left\{ \begin{array}{l} \langle \text{true}, \{m_3\} \rangle, \quad \langle \text{true}, m'_3 \rangle, \\ \langle z = \text{enc}(x), \{\text{ack}(\text{enc}(x)).\top, \text{if } z \neq \text{enc}(x) \text{ then } \text{ack}(z').\perp \text{ else } \text{ack}(z).\top\} \rangle, \\ \langle (z = \text{enc}(x)) \wedge (w = z) \wedge (w = \text{enc}(x)), \{\top\} \rangle, \\ \langle \neg(z = \text{enc}(x)), \{\text{if } z \neq \text{enc}(x) \text{ then } \text{ack}(z').\perp \text{ else } \text{ack}(z).\top\} \rangle, \quad \langle \neg(z = \text{enc}(x)), \{\perp\} \rangle \end{array} \right\}$$

In \mathcal{S}_2 and the ensuing discussion, we alleviate our presentation by simplifying the boolean conditions used, e.g., we simply write $(z = \text{enc}(x))$ in lieu of $(\text{true} \wedge (z = \text{enc}(x)))$. We highlight a few points.

First, consider the second constrained monitor-set in \mathcal{S}_2 , namely $\langle \text{true}, m'_3 \rangle$. Since the semantics of Figure 2 allows expression guards and quantified guards to transition with the *same* symbolic event (albeit with different conditions) we are able to consider the *resp.* continuations in unison for the event $\text{aut}\langle z \rangle$. Concretely, according to Definition 19(ii), for $\text{spa}(\langle \text{true}, m'_3 \rangle, \text{aut}\langle z \rangle, (z = \text{enc}(x)))$ generated by the *expression* guard weak transition of m'_3 , we need to ensure that the resulting monitor-set $\langle z = \text{enc}(x), \text{saft}(\{m'_3\}, \{(z = \text{enc}(x)), \text{true}\}, \text{aut}\langle z \rangle) \rangle$ (which evaluates to the third constrained monitor-set in \mathcal{S}_2) is also in the symbolic relation. At the same time, for $\text{spa}(\langle \text{true}, m'_3 \rangle, \text{aut}\langle z \rangle, \text{true})$ generated by the *quantified* guard weak transition of m'_3 , we need to ensure that two monitor-sets are in \mathcal{S}_2 , namely $\langle z = \text{enc}(x), \text{saft}(\{m'_3\}, \{(z = \text{enc}(x)), \text{true}\}, \text{aut}\langle z \rangle) \rangle$ (as before) but also the constrained monitor-set $\langle \neg(z = \text{enc}(x)), \text{saft}(\{m'_3\}, \{\neg(z = \text{enc}(x)), \text{true}\}, \text{aut}\langle z \rangle) \rangle$ (which evaluates to the fifth constrained monitor-set in \mathcal{S}_2).

The second point we highlight about \mathcal{S}_2 concerns its third constrained monitor-set. In particular, the left branch of the conditional term in this set, namely $\text{ack}(z').\perp$ in the term $\text{if } z \neq \text{enc}(x) \text{ then } \text{ack}(z').\perp \text{ else } \text{ack}(z).\top$, is not considered by our analysis since its condition, $z \neq \text{enc}(x)$, is *incompatible* with the constraining condition of the monitor-set, i.e., $\neg \text{saft}((z = \text{enc}(x)) \wedge (z \neq \text{enc}(x)))$.

Third, we also note how the condition aggregation mechanism for the consecutive symbolic events $\text{aut}\langle z \rangle$ and $\text{ack}\langle w \rangle$ — transferring us from the second constrained monitor-set, $\langle \text{true}, m'_3 \rangle$, to the fourth, $\langle (z = \text{enc}(x)) \wedge (w = z) \wedge (w = \text{enc}(x)), \{\top\} \rangle$, via the third constrained monitor-set in \mathcal{S}_2 — enables us to symbolically relate the *expression* guards in m_3 , which impose a condition such as $(z = \text{enc}(x))$ upon transition, with the *quantified* guard that imposes the same condition *after* the transition (by means of a conditional branch in its continuation). We leave it up to the interested reader to check that the remaining monitor-sets in \mathcal{S}_2 satisfy the conditions required by Definition 19. ◀

6 On Automating Symbolic Controllability

Despite its merits, a direct implementation of the symbolic controllability from Definition 19 still would *not* perform well for certain recursive monitor descriptions, as shown in the following example.

► **Example 23.** Recall monitor m_4 from Example 2. To show that it is controllable, we need to exhibit a symbolic relation that includes $\langle \text{true}, \{m_4\} \rangle$. For some fresh variable x where $\text{frsh}(\text{fv}(\langle \text{true}, \{m_4\} \rangle)) = x$, since the judgement $\text{spa}(\langle \text{true}, \{m_4\} \rangle, \text{in}\langle x \rangle, \text{true})$ holds, this relation needs to include the ensuing monitor-set $\langle \text{true}, \{m'_4\} \rangle$ as well, where $m'_4 \triangleq \text{if } x = 80 \text{ then } \text{out}\langle 81 \rangle.\top \text{ else } \text{out}\langle x \rangle.m_4$. In turn, since $\text{spa}(\langle \text{true}, \{m'_4\} \rangle, \text{out}\langle y \rangle, (\neg(x = 80) \wedge y = x))$ (where $\text{frsh}(\text{fv}(\langle b, \{m'_4\} \rangle)) = y$), the symbolic relation must also contain $\langle \neg(x = 80) \wedge y = x, \{m_4\} \rangle$. We thus reach the original monitor set $\{m_4\}$ but with a stronger condition, namely

$\neg(x=80) \wedge y=x$. By extension of this reasoning, it is not hard to see that the symbolic relation required by Definition 19 needs to be infinitely large. ◀

The problem exhibited by Example 23 is that the condition aggregating mechanism of Definition 19 does not specify any means for consolidating the boolean condition b constraining a monitor set M in $\langle b, M \rangle$, i.e., a form of garbage collection of redundant conditions. For instance, in the constrained monitor-set $\langle \neg(x=80) \wedge y=x, \{m_4\} \rangle$ of Example 23, the condition $(\neg(x=80) \wedge y=x)$ plays *no effective role* in constraining the free variables in $\{m_4\}$, of which there are none. We therefore optimise Definition 19 in a sound (and complete) manner by taking into consideration boolean sub-conditions that can be isolated and discarded. This leads to an improved automated analysis for consistently-detecting monitors.

► **Definition 24** (Optimised Symbolic Controllability). The consolidation of a boolean expression b wrt. a variable set V , denoted as $\mathbf{cns}(b, V)$, is defined as:

$$\mathbf{cns}(b, V) \stackrel{\text{def}}{=} b_1 \text{ whenever } \mathbf{prt}(b, V) = \langle b_1, b_2 \rangle \text{ for some } b_2$$

where the boolean expression partitioning operation $\mathbf{prt}(b, V)$ is defined as:

$$\mathbf{prt}(b, V) \stackrel{\text{def}}{=} \begin{cases} \langle b_1, b_2 \rangle & \text{if } \mathbf{sat}(b) \text{ and } b = b_1 \wedge b_2 \text{ and } (\mathbf{fv}(b_1) \subseteq V) \text{ and } (V \cap \mathbf{fv}(b_2) = \emptyset) \\ \langle b, \text{true} \rangle & \text{otherwise} \end{cases}$$

Let the *optimised symbolic reachability* from $\langle b, M \rangle$ for θ and c , $\mathbf{osoft}(\langle b, M \rangle, \theta, c)$, be defined as:

$$\mathbf{osoft}(\langle b, M \rangle, \theta, c) \stackrel{\text{def}}{=} \left\{ \langle \mathbf{cns}(\wedge B, V), \mathbf{soft}(M, B, \theta) \rangle \mid \begin{array}{l} B \in \mathbf{sc}(b \wedge c, \mathbf{rc}(M, \theta)) \text{ and } \mathbf{sat}(\wedge B) \\ \text{and } V = \mathbf{fv}(\mathbf{soft}(M, B, \theta)) \end{array} \right\}$$

A relation $\mathcal{S} \subseteq (\text{BEXP} \times \mathcal{P}(\text{MON}))$ is called *optimised symbolically-controllable* iff for all $\langle b, M \rangle \in \mathcal{S}$:

1. $\mathbf{spr}(\langle b, M \rangle, w)$ and $w \in \{\top, \perp\}$ implies $M = \{w\}$;
2. $\mathbf{spa}(\langle b, M \rangle, l(x), c)$ where $\mathbf{frsh}(\mathbf{fv}(\langle b, M \rangle)) = x$ implies $\mathbf{osoft}(\langle b, M \rangle, l(x), c) \subseteq \mathcal{S}$.

The *largest* optimised symbolically-controllable relation is denoted by $\mathcal{C}_{\text{sym}}^{\text{opt}}$. A (closed) monitor m is said to be optimised symbolically-controllable iff $\langle \text{true}, \{m\} \rangle \in \mathcal{C}_{\text{sym}}^{\text{opt}}$. ◀

We highlight the salient points from Definition 24. First, boolean consolidation in a constrained monitor-set, $\langle \mathbf{cns}(b), M \rangle$, should *not* change the set of concrete monitor sets represented by $\langle b, M \rangle$ and, for this reason, we cannot consolidate *unsatisfiable* boolean conditions. For instance, even if $x \notin \mathbf{fv}(M)$, it is still unsound to optimise $\langle \text{true} \wedge (x \neq x), M \rangle$ to $\langle \text{true}, M \rangle$ based on the fact that $\neg \mathbf{sat}(x \neq x)$. Concretely, from Equation 4 of Section 5 we know that $\langle \text{true} \wedge x \neq x, M \rangle$ denotes the empty set of monitor-sets, \emptyset , whereas $\langle \text{true}, M \rangle$ represents the set $\{m\rho \mid \llbracket \text{true} \rrbracket \rho = \text{true} \text{ and } m \in M\}$. Second, consolidation should ideally filter out as much redundant constraints as possible, e.g., in $\langle b_1 \wedge b_2, M \rangle$ we should remove b_2 whenever $\mathbf{fv}(b_2) \cap \mathbf{fv}(M) = \emptyset$. In Definition 24 we require the strongest possible condition for the residual condition b_1 in $\langle b_1 \wedge b_2, M \rangle$, i.e., $\mathbf{fv}(b_1) \subseteq \mathbf{fv}(M)$, which indirectly implies that the *resp.* condition variables are partitioned $\mathbf{fv}(b_1) \cap \mathbf{fv}(b_2) = \emptyset$. This partitioning is crucial for a sound consolidation, e.g., in $\langle (\neg(x=80) \wedge y=x), M \rangle$, it is unsound to just remove the subcondition $\neg(x=80)$ when $x \notin \mathbf{fv}(M)$ and $y \in \mathbf{fv}(M)$. Although $\mathbf{prt}(b, V)$ can be refined further (while still observing core requirements for soundness such as variable condition partitioning), in Definition 24 we opted for a less elaborate condition that suffices our exposition. Third, we highlight the fact that the conditions specifying $\mathbf{cns}(b)$ in Definition 24 yield a unique consolidated condition up to semantic equivalence meaning that, in an implementation of the framework, this can be defined as a function.

► **Theorem 25** (Optimised Controllability). $\langle \text{true}, \{m\} \rangle \in \mathcal{C}_{sym}^{opt}$ iff $\langle \text{true}, \{m\} \rangle \in \mathcal{C}_{sym}$ ◀

► **Example 26.** As a result of Theorems 13, 21 and 25, we can show that m_4 of Example 2 is consistently-detecting by exhibiting the optimised symbolic relation \mathcal{S}_3 below (m'_4 is the monitor defined earlier in Example 23). For expository purposes, we show how the consolidated boolean expressions are calculated. In particular, the first constrained monitor-set in \mathcal{S}_3 denotes both the starting pair $\langle \text{true}, \{m_4\} \rangle$, but also the pair $\langle \mathbf{cns}(x = 80 \wedge y = 81, \mathbf{fv}(\{m_4\})), \{m_4\} \rangle$.

$$\mathcal{S}_3 = \left\{ \left\langle \underbrace{\text{true}}_{\mathbf{cns}(\neg(x=80) \wedge y=x, \emptyset)}, \{m_4\} \right\rangle, \left\langle \underbrace{\text{true}}_{\mathbf{cns}(x=80, \emptyset)}, \{\perp\} \right\rangle, \langle \text{true}, \{m'_4\} \rangle, \left\langle \underbrace{\text{true}}_{\mathbf{cns}(x=80 \vee y=81, \emptyset)}, \{\top\} \right\rangle \right\}$$

The interested reader may check that the conditions of Definition 24 are satisfied by \mathcal{S}_3 . ◀

Using standard techniques [42, 3], an algorithm constructing symbolically-controllable relations from Definition 24 can be extracted more easily. Moreover, the completeness aspect in Theorems 13, 21, and 25 should enable such an automation to infer a *counter-example* (system and trace) from failed attempts, thereby explaining why a monitor is *not* consistently-detecting.

► **Example 27.** Recall monitor m_2 from Example 1. Assuming the shorthand abbreviations $m'_2 \triangleq \text{aut}\langle y \rangle. \text{ack}\langle y \rangle. \top + \text{aut}\langle z \rangle. \text{if } z \neq y \text{ then } \text{ack}\langle z' \rangle. \perp$ and $M = \{\text{if } z \neq \text{enc}(x) \text{ then } \text{ack}\langle z' \rangle. \perp, \text{ack}\langle \text{enc}(x) \rangle. \top\}$, compiling a relation satisfying Definition 24 fails because it needs to include:

$$\begin{array}{ll} \langle \text{true}, \{\text{let } y = \text{enc}(x) \text{ in } m'_2\} \rangle & \text{since } \mathbf{spa}(\langle \text{true}, \{m_2\} \rangle, \text{chl}\langle x \rangle, \text{true}) \\ \langle z = \text{enc}(x), M \rangle & \text{since } \mathbf{spa}(\langle \text{true}, \{\text{let } y = \text{enc}(x) \text{ in } m'_2\} \rangle, \text{aut}\langle z \rangle, z = \text{enc}(x)) \\ \langle \text{true}, \{\top, \mathbf{0}\} \rangle & \text{since } \mathbf{spa}(\langle z = \text{enc}(x), M \rangle, \text{ack}\langle u \rangle, (u = \text{enc}(x))) \end{array}$$

The final pair $\langle \text{true}, \{\top, \mathbf{0}\} \rangle$ violates Definition 24(i) and, via the symbolic events on right-hand column conatining the $\mathbf{spa}(-)$ assertions that lead to it, one can construct the counter-example inducing the inconsistent detection, *i.e.*, a system s producing the trace $\text{chl}\langle v \rangle \cdot \text{aut}\langle u \rangle \cdot \text{ack}\langle u \rangle$ where $u = \text{enc}(v)$. ◀

7 Conclusion

Monitors should provide guarantees that they will operate correctly when instrumented with *any* system [35]. At the same time, for this requirement to be scalable, the corresponding correctness analysis that determines it must also be *compositional*: monitors should be verified separately, independent of the systems they may be instrumented with. The fact that monitors tend to be considerably smaller (in size and complexity) than the systems they observe further justifies this point. This paper provides two definitions that formalise deterministic monitoring behaviour, Definition 6 (consistently-detecting monitors) and Definition 11 (controllable monitors), that address these seemingly conflicting concerns; it also shows that the two definitions coincide, Theorem 13. In addition, the paper also studies alternative definitions for controllability, Definitions 19 and 24, that enable the implementation of sound and complete symbolic analyses, Theorems 21 and 25.

Our methods provide a systematic way for factoring out auxiliary reasoning on data from the analysis relating to the branching structure of the monitors; the former kind of reasoning can be determined by calling on an independent satisfiability solver. In fact, for the specific case of our expository monitor language in Figure 1, one can show that our methods yield *finite* symbolic transition graphs, making the latter reasoning decidable modulo the expression and boolean language used. The results obtained in this work should also be general enough

to be applicable to other monitoring systems. For instance, Definitions 6, 11, 19, and 24 are independent of the kind of systems monitored, the syntax of the monitor language, and the event value domains and expression languages used. Instead, they are defined in terms of generic characteristics such as their LTS semantics. As a result, extending the monitor language with constructs such as parallel composition would not affect the existing framework. The instrumentation relation one adopts for composing monitors with systems necessarily affects the compositional properties and the correspondence between the respective definitions. However, these changes do not impinge on the general structure of our definitions and should be local to the detection condition in the *resp.* controllability definitions, namely Definition 11(*i*), 19(*i*) and 24(*i*), and the reachability-set definitions **aft**(-), **saft**(-) and **osaft**(-) of Definitions 10, 18, and 24.

Future work. We will further investigate the implementability aspects of our analysis, possibly as an extension to existing model-checking tools. This may raise further issues and adjustments to our definitions *e.g.*, it may be more efficient to batch the satisfiability checks in Definition 24. We plan to apply this to existing transition-based monitor specifications such as [5, 46] and validate its feasibility as an automated specification assistant.

Related Work. The need for determinising monitor syntheses from logical specifications is frequently discussed in the literature [7, 23]. In [1], the authors employ a trace-based definition of deterministic monitors that takes into consideration verdicts (similar to our definition for consistent detections of Definition 6 but without considering universal quantification over system instrumentations) and establish complexity bounds for determinising monitors *wrt.* this definition. Set-simulations, which are related to our monitor-sets, have been used as a proof technique for testing preorders in [15] but do not consider symbolic analyses. Acceptor ambiguity [29] is closely related to our notion of consistent detection with respect to the three outcomes of acceptance, rejection and withholding, as specified in Definition 6. Crucially, however, our definition universally quantifies over all possible system compositions. Subsequently, the main endeavour of our work was that of developing sound and complete compositional techniques to alleviate the analysis for consistent detection; we are unaware of any compositional or coinductive techniques used for determining acceptor ambiguity. Symbolic LTSs were studied extensively for value-passing CCS in [27, 28], but their use in controllability for reasoning about consistent monitor detections is, to our knowledge, novel. The particular setting where it is used, namely the instrumentation composition relation and the use of monitor-sets, also require new technical machinery, such as that of Definition 18. Our definition of controllability, Definition 11, is related to viability (usability) for clients in contract compliance [39] and must-testing [9]. Particularly, in the case of compliance, viability is defined coinductively and is satisfied whenever there *exists* a server that can engage with the client so as to lead it to success whenever interaction terminates. Apart from the universal quantifications over systems (viability existentially quantifies over servers), our work differs from [39, 9] *wrt.* the treatment of verdicts considered, the composition relation used (*i.e.*, instrumentation), and the development of a symbolic analysis for handling of action/event data.

Acknowledgements. The author acknowledges the Dagstuhl seminar 17051 and would like to thank Luca Aceto, Antonis Achilleos, Duncan Attard, Giovanni Bernardi, Ian Cassar, Anna Ingólfssdóttir, Giles Reger and anonymous reviewers for their help and suggestions.

References

- 1 Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir, and Sævar Örn Kjartansson. On the Complexity of Determinizing Monitors. In *CIAA*, pages 1–13, 2017. doi:10.1007/978-3-319-60134-2_1.
- 2 Luca Aceto, Anna Ingólfssdóttir, Kim Guldstrand Larsen, and Jiri Srba. *Reactive Systems: Modelling, Specification and Verification*. Cambridge Univ. Press, 2007.
- 3 Luca Aceto, Anna Ingólfssdóttir, and Jiri Srba. *Advanced Topics in Bisimulation and Coinduction*, chapter The Algorithmics of Bisimilarity. Cambridge Univ. Press, 2011.
- 4 Cyril Allauzen, Mehryar Mohri, and Ashish Rastogi. General algorithms for testing the ambiguity of finite automata. In *DTL*, pages 108–120, 2008. doi:10.1007/978-3-540-85780-8_8.
- 5 Howard Barringer, Yliès Falcone, Klaus Havelund, Giles Reger, and David E. Rydeheard. Quantified Event Automata: Towards Expressive and Efficient Runtime Monitors. In *FM*, pages 68–84, 2012. doi:10.1007/978-3-642-32759-9_9.
- 6 David Basin, Felix Klaedtke, Srdjan Marinovic, and Eugen Zălinescu. Monitoring compliance policies over incomplete and disagreeing logs. In *RV*, pages 151–167, 2013. doi:10.1007/978-3-642-35632-2_17.
- 7 Andreas Bauer, Martin Leucker, and Christian Schallhart. Runtime verification for LTL and TLTL. *TOSEM*, 20(4):14, 2011.
- 8 Shay Berkovich, Borzoo Bonakdarpour, and Sebastian Fischmeister. Runtime verification with minimal intrusion through parallelism. *Formal Methods in System Design*, 46(3):317–348, 2015. doi:10.1007/s10703-015-0226-3.
- 9 Giovanni Bernardi and Adrian Francalanza. Full-abstraction for must testing preorders. In *COORDINATION*, pages 237–255, 2017. doi:10.1007/978-3-319-59746-1_13.
- 10 Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. *TCS*, 669:33–58, 2017. doi:10.1016/j.tcs.2017.02.009.
- 11 Borzoo Bonakdarpour, Pierre Fraigniaud, Sergio Rajsbaum, David A. Rosenblueth, and Corentin Travers. Decentralized asynchronous crash-resilient runtime verification. In *CONCUR*, pages 16:1–16:15, 2016. doi:10.4230/LIPIcs.CONCUR.2016.16.
- 12 Ian Cassar and Adrian Francalanza. Runtime Adaptation for Actor Systems. In *RV*, volume 9333, pages 38–54. Springer, 2015.
- 13 Ian Cassar and Adrian Francalanza. On Implementing a Monitor-Oriented Programming Framework for Actor Systems. In *iFM*, pages 176–192, 2016.
- 14 Feng Chen and Grigore Roşu. MOP: An Efficient and Generic Runtime Verification Framework. In *OOPSLA*, pages 569–588, 2007. doi:10.1145/1297027.1297069.
- 15 Rance Cleaveland and Matthew Hennessy. Testing equivalence as a bisimulation equivalence. *FACS*, 5(1):1–20, 1993.
- 16 Christian Colombo, Adrian Francalanza, Ruth Mizzi, and Gordon J. Pace. polylarva: Runtime verification with configurable resource-aware monitoring boundaries. In *SEFM*, pages 218–232, 2012.
- 17 Marcelo d’Amorim and Grigore Roşu. Efficient monitoring of ω -languages. In *CAV*, pages 364 – 378, 2005.
- 18 Søren Debois, Thomas Hildebrandt, and Tijs Slaats. Safety, liveness and run-time refinement for modular process-aware systems with dynamic sub processes. In *FM*, pages 143–160, 2015. doi:10.1007/978-3-319-19249-9_10.
- 19 Normann Decker and Daniel Thoma. On freeze LTL with ordered attributes. In *FOSSACS*, pages 269–284, 2016. doi:10.1007/978-3-662-49630-5_16.
- 20 John Dorsey. *Continuous and Discrete Control Systems: Modeling, Identification, Design, and Implementation*. McGraw-Hill, 2001.

- 21 Pierre Fraigniaud, Sergio Rajsbaum, and Corentin Travers. On the number of opinions needed for fault-tolerant run-time monitoring in distributed systems. In *RV*, pages 92–107, 2014. doi:10.1007/978-3-319-11164-3_9.
- 22 Adrian Francalanza. A Theory of Monitors. In *FoSSaCS*, pages 145–161, 2016.
- 23 Adrian Francalanza, Luca Aceto, and Anna Ingolfsdottir. Monitorability for the Hennessy–Milner logic with recursion. *FMSD*, pages 1–30, 2017. doi:10.1007/s10703-017-0273-z.
- 24 Adrian Francalanza and Aldrin Seychell. Synthesising Correct concurrent Runtime Monitors. *FMSD*, 46(3):226–261, 2015. doi:10.1007/s10703-014-0217-9.
- 25 Jan Friso Groote and M.P.A. Sellink. Confluence for process verification. *TCS*, 170(1):47–81, 1996. doi:10.1016/S0304-3975(96)80702-X.
- 26 Yuqin He, Xiangping Chen, and Ge Lin. Composition of monitoring components for on-demand construction of runtime model based on model synthesis. In *Internetware*, pages 1–5, 2013. doi:10.1145/2532443.2532472.
- 27 Matthew Hennessy and Anna Ingolfsdottir. A Theory of Communicating Processes with Value Passing. *Information and Computation*, 107(2):202–236, 1993. doi:10.1006/inco.1993.1067.
- 28 Matthew Hennessy and Huimin Lin. Symbolic bisimulations. *TCS*, 138(2):353–389, 1995. doi:10.1016/0304-3975(94)00172-F.
- 29 Oscar H. Ibarra and Bala Ravikumar. On Sparseness, Ambiguity and other decision problems for Acceptors and Transducers. In *STACS*, pages 171–179, 1986. doi:10.1007/3-540-16078-7_74.
- 30 Limin Jia, Hannah Gommerstadt, and Frank Pfenning. Monitors and blame assignment for higher-order session types. In *POPL*, pages 582–594, 2016. doi:10.1145/2837614.2837662.
- 31 Jerzy Klamka. *Control System, Robotics and Automation*, volume 7, chapter System Characteristics: Stability, Controllability, Observability. EOLLS, 2009.
- 32 John Klein and Ian Gorton. Runtime Performance Challenges in Bigdata Systems. In *WOSP*, pages 17–22, 2015. doi:10.1145/2693561.2693563.
- 33 D. König. Über eine Schlussweise aus dem Endlichen ins Unendliche. *Acta Litt. ac. sci. Szeged*, 3, 1927.
- 34 Dexter Kozen. Results on the propositional μ -calculus. *TCS*, 27:333–354, 1983.
- 35 Jonathan Laurent, Alwyn Goodloe, and Lee Pike. Assuring the Guardians. In *RV*, pages 87–101, 2015.
- 36 Jay Ligatti, Lujo Bauer, and David Walker. Edit automata: enforcement mechanisms for run-time security policies. *Int. J. Inf. Secur.*, 4(1-2):2–16, 2005. doi:10.1007/s10207-004-0046-8.
- 37 Qingzhou Luo and Grigore Roşu. EnforceMOP: A Runtime Property Enforcement System for Multithreaded Programs. In *ISSTA*, New York, NY, USA, 2013. ACM. doi:10.1145/2483760.2483766.
- 38 Patrick O’Neil Meredith, Dongyun Jin, Dennis Griffith, Feng Chen, and Grigore Roşu. An overview of the MOP runtime verification framework. *STTT*, 14(3):249–289, 2012. doi:10.1007/s10009-011-0198-6.
- 39 Luca Padovani. Contract-based discovery of web services modulo simple orchestrators. *TCS*, 411(37), 2010.
- 40 Anna Philippou and David Walker. On confluence in the π -calculus. In *ICALP*, pages 314–324, 1997. doi:10.1007/3-540-63165-8_188.
- 41 Giles Reger, Helena Cuenca Cruz, and David E. Rydeheard. MarQ: Monitoring at Runtime with QEA. In *TACAS*, pages 596–610, 2015. doi:10.1007/978-3-662-46681-0_55.
- 42 Davide Sangiorgi. *An introduction to Bisimulation and Coinduction*. Cambridge University Press, 2012.

- 43 Fred B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*, 3(1):30–50, 2000. doi:10.1145/353323.353382.
- 44 Michael Sipser. *Introduction to the Theory of Computation*. Cengage Learning, 3rd edition, 2012.
- 45 Moshe Y. Vardi and Pierre Wolper. Reasoning about infinite computations. *Inf. & Comp.*, 115(1):1–37, 1994. doi:10.1006/inco.1994.1092.
- 46 Yoriyuki Yamagata, Cyrille Artho, Masami Hagiya, Jun Inoue, Lei Ma, Yoshinori Tanabe, and Mitsuharu Yamamoto. Runtime monitoring for concurrent systems. In *RV*, pages 386–403, 2016. doi:10.1007/978-3-319-46982-9_24.