

Comparing Controlled System Synthesis and Suppression Enforcement

Luca Aceto · Ian Cassar · Adrian Francalanza · Anna Ingólfssdóttir

Received: date / Accepted: date

Abstract Runtime enforcement and control system synthesis are two verification techniques that automate the process of transforming an erroneous system into a valid one. As both techniques can modify the behaviour of a system to prevent erroneous executions, they are both ideal for ensuring safety. In this paper, we investigate the interplay between these two techniques and identify control system synthesis as being the static counterpart to suppression-based runtime enforcement, in the context of safety properties.

This work was partly supported by the projects “TheoFoMon: Theoretical Foundations for Monitorability” (nr.163406-051) and “Developing Theoretical Foundations for Runtime Enforcement” (nr.184776-051) of the Icelandic Research Fund, by the EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement nr. 778233, and by the Endeavour Scholarship Scheme (Malta), part-financed by the European Social Fund (ESF) - Operational Programme II Cohesion Policy 2014-2020.

L. Aceto
Gran Sasso Science Institute, L’Aquila, Italy
Department of Computer Science, Reykjavík University, Iceland
E-mail: luca@ru.is

I. Cassar
Department of Computer Science, University of Malta, Malta
Department of Computer Science, Reykjavík University, Iceland
E-mail: ian.cassar.10@um.edu.mt

A. Francalanza
Department of Computer Science, University of Malta, Malta
E-mail: adrian.francalanza@um.edu.mt

A. Ingólfssdóttir
Department of Computer Science, Reykjavík University, Iceland
E-mail: annai@ru.is

1 Introduction

Our increasing reliance on software systems is raising the demand for ensuring their reliability and correctness. Several verification techniques help facilitate this task by automating the process of deducing whether the system under scrutiny (SuS) satisfies a predefined set of correctness properties. Properties are either verified pre-deployment (statically), using techniques such as model checking (MC) [3, 13], or post-deployment (dynamically), as per runtime verification (RV) [12, 21, 28]. In both cases, any error discovered during the verification serves as guidance for identifying the invalid parts of the system that require adjustment.

Other post-deployment techniques, such as *runtime enforcement* (RE), additionally attempt to automatically transform the invalid system into a valid one. Runtime enforcement [6, 16, 27, 29] adopts an intrusive monitoring approach by which every observable action executed by the SuS is scrutinized and modified as necessary by a monitor at runtime. Monitors in RE may be described in various ways, such as: transducers [6, 9, 33], shields [27] and security automata [18, 29, 35]. They may opt to *replace* the invalid actions by valid ones, or completely *suppress* them, thus rendering them immaterial to the environment interacting with the SuS; in certain cases, monitors can even *insert* actions that may directly affect the environment. Different enforcement strategies are applied depending on the property that needs to be enforced.

A great deal of effort [7, 14, 23, 24, 26] has been devoted to studying the interplay between static and dynamic techniques, particularly to understand how the two can be used in unison to minimise their respective weaknesses. It is well established that runtime verification is the *dynamic counterpart* of model checking,

which means that a subset of the properties verifiable using MC can also be verified dynamically via RV. In fact, multi-pronged verification approaches often use RV in conjunction with MC. For instance, system verification based on MC is often carried out with respect to a model of the environment in which the studied system operates. This makes the sole use of model-checking techniques problematic in settings, such as mobile robotics, where the precise conditions in which systems operate are often known only at runtime and may change over time. Here RV can be used to check post-deployment the environmental conditions used to validate systems at design time [14]. Sometimes, MC is used to statically verify the parts of the SuS which cannot be verified dynamically (e.g., due to inaccessible code or high impact on the performance of the SuS), while RV is then used to verify other parts dynamically in order to minimise the state explosion problem inherent to MC.

A natural question to ask is which technique can be considered as the *static counterpart* to runtime enforcement, i.e., a technique that can statically achieve the same (or equivalent) results as per runtime enforcement. Identifying such a technique is quite desirable as it would allow for properties to be enforced statically when dynamic verification is not ideal, e.g., when the monitor’s runtime overheads are infeasible, and vice versa e.g., to mitigate state explosions during static analysis. Such a technique therefore enables the possibility of adopting a multi-pronged enforcement approach that is similar to the one used for verification with RV and MC, but from an enforcement perspective. One promising static technique that has several aspects in common with runtime enforcement is *controlled system synthesis* (CSS) [10, 15, 25, 31]. This approach analyses the state space of the SuS and reformulates it pre-deployment to *remove* the system’s ability of executing erroneous behaviour. As a result, a restricted (yet valid) version of the SuS is produced; this is known as a *controlled system*.

The primary aim of both RE and CSS is to force the resulting monitored/controlled system to adhere to the respective property – this is known as *soundness* in RE and *validity* in CSS. Further guarantees are also generally required to ensure minimal disruption to valid systems – this is ensured via *transparency* in RE and *maximal permissiveness* in CSS. As both techniques may adjust systems by omitting their invalid behaviours, they are ideal for ensuring *safety*. These observations, along with other commonalities, hint at the existence of a relationship between runtime enforcement and controlled system synthesis, in the context of safety properties.

In this paper we conduct a preliminary investigation on the interplay between the above mentioned two techniques with the aim of establishing a static counterpart

for runtime enforcement. We intend to identify a set of properties that can be enforced either dynamically, via runtime enforcement, or statically via controlled system synthesis. In this first attempt, we however limit ourselves to study this relationship in the context of *safety properties*. As a vehicle for this comparison, we choose the recent work on CSS by van Hulst *et al.* [25], and compare it to our previous work, presented in [6], on enforcing safety properties via action suppressions. We chose these two bodies of work as they are accurate representations of the two techniques. Moreover, they share a number of commonalities including their choice of specification language, modelling of systems, *etc.* To further simplify our comparison, we formulate both techniques in a core common setting and show that there are subtle differences between them even in that scenario. Specifically, we identify a common core within the work presented in [6, 25] by:

- working with respect to the Safe Hennessy Milner Logic with invariance (SHML_{inv}), that is, the *intersection* of the logics used by both works, namely, the Safe Hennessy Milner Logic with recursion (SHML) in [6] and the Hennessy Milner Logic with invariance and reachability ($\text{HML}_{\text{inv}}^{\text{reach}}$) in [25],
- removing constructs and aspects that are supported by one technique and not by the other, and by
- taking into account the assumptions considered in both bodies of work.

To our knowledge, no one has yet attempted to identify a static counterpart to RE, and an insightful comparison of RE and CSS has not yet been conducted. As part of our investigation we offer the following contributions:

- (i) We prove that the monitored system obtained from instrumenting a suppression monitor derived from a formula, and the controlled version of the same system (by the same formula) are *trace (language) equivalent*, that is, they can execute the same set of traces, Theorem 2.
- (ii) When restricted to safety properties, controlled system synthesis is the *static counterpart* (Definition 3) to runtime enforcement, Theorem 3.
- (iii) In spite of (i) and (ii), both of the obtained systems need not be observationally equivalent, Theorem 4.

Although (i) suffices to deduce (ii) since it is well known that trace equivalent systems satisfy the exact same set of safety properties, Theorem 1, (iii) entails that a very powerful external observer can, at least in principle, tell the difference between these two resultant systems [1].

Structure of the paper. This article is an extended version of [8]; it includes improved and more expanded

explanations and examples, along with the complete proofs for our theorems.

The rest of the paper is structured as follows. Section 2 provides the necessary preliminary material describing how we model systems as labelled transition systems and properties via the chosen logic. In Section 3 we give an overview of the simplified versions of the enforcement model presented in [6] and the controlled system synthesis rules of [25]. In Section 4 we then present our first set of contributions consisting of a mapping function that derives enforcement monitors from logic formulas, and the proof that the obtained monitored and controlled versions of a given system are *trace equivalent*. Section 5 then presents a deeper comparison of the differences and similarities between the two models, followed by our second contribution which *disproves* the observational equivalence of the two techniques. This allows us to establish that controlled system synthesis is the *static counterpart* to enforcement when it comes to safety properties. Section 6 overviews related work, and Section 7 concludes.

2 Preliminaries

The Model: We assume systems described as *labelled transition systems* (LTSs), which are triples $\langle \text{SYS}, (\text{ACT} \cup \{\tau\}), \rightarrow \rangle$ defining a set of *system states*, $s, r, q \in \text{SYS}$, a finite set of *observable actions*, $\alpha, \beta \in \text{ACT}$, and a distinguished silent action $\tau \notin \text{ACT}$, along with a *transition relation*, $\rightarrow \subseteq (\text{SYS} \times \text{ACT} \cup \{\tau\} \times \text{SYS})$. We let $\mu \in \text{ACT} \cup \{\tau\}$ and write $s \xrightarrow{\mu} r$ in lieu of $(s, \mu, r) \in \rightarrow$. We use $s \xrightarrow{\alpha} r$ to denote weak transitions representing $s \xrightarrow{(\tau)^*} r$ and refer to r as an α -derivative of s . Traces $t, u \in \text{ACT}^*$ range over (finite) sequences of observable actions, and we write $s \xrightarrow{t} r$ for 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 $n = 0$, t is the empty trace ε and $s \xrightarrow{\varepsilon} r$ means $s \xrightarrow{(\tau)^*} r$. For each $\mu \in \text{ACT} \cup \{\tau\}$, the notation $\hat{\mu}$ stands for ε if $\mu = \tau$ and for μ otherwise. We write $\text{traces}(s)$ for the set of traces executable from system state s , that is, $t \in \text{traces}(s)$ iff $s \xrightarrow{t} r$ for some r . We use the syntax of the regular fragment of CCS [30] to concisely describe LTSs in our examples. We also assume the classic notions for *trace (language) equivalence* and *observational equivalence*, that is, weak bisimilarity [30, 34].

Definition 1 (Trace Equivalence) Two LTS system states s and r are *trace equivalent* iff they produce the same set of traces, i.e., $\text{traces}(s) = \text{traces}(r)$. \square

Definition 2 (Observational Equivalence) A relation \mathcal{R} over a set of system states is a *weak bisimulation*

Syntax

$$\begin{array}{ll} \varphi, \psi \in \text{sHML} ::= \text{tt} \text{ (truth)} & | \text{ff} \text{ (falsehood)} \\ & | \varphi \wedge \psi \text{ (conjunction)} & | [\alpha]\varphi \text{ (necessity)} \\ & | \max X.\varphi \text{ (greatest fp.)} & | X \text{ (fp. variable)} \end{array}$$

Semantics

$$\begin{aligned} \llbracket \text{tt}, \rho \rrbracket &\stackrel{\text{def}}{=} \text{SYS} & \llbracket \text{ff}, \rho \rrbracket &\stackrel{\text{def}}{=} \emptyset & \llbracket X, \rho \rrbracket &\stackrel{\text{def}}{=} \rho(X) \\ \llbracket [\alpha]\varphi, \rho \rrbracket &\stackrel{\text{def}}{=} \{s \mid \forall r \cdot s \xrightarrow{\alpha} r \text{ then } r \in \llbracket \varphi, \rho \rrbracket\} \\ \llbracket \varphi \wedge \psi, \rho \rrbracket &\stackrel{\text{def}}{=} \llbracket \varphi, \rho \rrbracket \cap \llbracket \psi, \rho \rrbracket \\ \llbracket \max X.\varphi, \rho \rrbracket &\stackrel{\text{def}}{=} \bigcup \{S \mid S \subseteq \llbracket \varphi, \rho[X \mapsto S] \rrbracket\} \end{aligned}$$

We also encode $\Box\varphi$ as $\max X.\varphi \wedge \bigwedge_{\beta \in \text{ACT}} [\beta]X$ where X is a fresh variable.

Fig. 1 The syntax and semantics for sHML.

iff whenever $(s, r) \in \mathcal{R}$ the following transfer properties are satisfied, for every action μ :

- $s \xrightarrow{\mu} s'$ implies there exists a transition $r \xrightarrow{\hat{\mu}} r'$ such that $(s', r') \in \mathcal{R}$; and
- $r \xrightarrow{\mu} r'$ implies there exists a transition $s \xrightarrow{\hat{\mu}} s'$ such that $(s', r') \in \mathcal{R}$.

Two system states s and r are *observationally equivalent*, denoted by $s \approx r$, iff there exists a weak bisimulation that relates them. \square

The Logic: The safety logic sHML [2, 3] is defined as the set of formulas generated by the grammar of Figure 1. It assumes a countably infinite set of logical variables $X, Y \in \text{LVAR}$ and provides the standard constructs of truth, tt, falsehood, ff, and conjunctions, $\varphi \wedge \psi$. As a shorthand, we occasionally denote conjunctions as $\bigwedge_{i \in I} \varphi_i$, where I is a finite set of indices, and when $I = \emptyset$, $\bigwedge_{i \in \emptyset} \varphi_i$ stands for tt. The logic is also equipped with the *necessity (universal) modality*, $[\alpha]\varphi$, and allows for defining recursive properties using greatest fixpoints, $\max X.\varphi$, which bind free occurrences of X in φ . We additionally encode the *invariance operator*, $\Box\varphi$, requiring φ to be satisfied by every reachable system state, as the recursive property, $\max X.\varphi \wedge \bigwedge_{\beta \in \text{ACT}} [\beta]X$, where X is not free in φ .

Formulas in sHML are interpreted over the system powerset domain where $S \in \mathcal{P}(\text{SYS})$. The semantic definition of Figure 1, $\llbracket \varphi, \rho \rrbracket$, 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. In that definition, $\rho' = \rho[X \mapsto S]$ denotes the valuation where $\rho'(X) = S$ and $\rho'(Y) = \rho(Y)$ for all other $Y \neq X$. We assume *closed* formulas, i.e., without free logical variables, and write $\llbracket \varphi \rrbracket$ in lieu of $\llbracket \varphi, \rho \rrbracket$ since the interpretation of a closed formula φ is independent of the valuation ρ . A system (state) s *satisfies* formula φ whenever $s \in \llbracket \varphi \rrbracket$.

$\varphi, \psi \in \text{sHML}_{\text{inv}} ::= \text{tt} \mid \text{ff} \mid \varphi \wedge \psi \mid [\alpha]\varphi \mid \square\varphi$

Fig. 2 The syntax for sHML_{inv} .

Example 1 Consider two systems (a good system, $s_{\mathbf{g}}$, and a bad one, $s_{\mathbf{b}}$) implementing a server that repeatedly accepts *requests*, *answers* them in response, and *logs* the serviced request. It also terminates upon accepting a *close* request. Whereas $s_{\mathbf{g}}$ outputs a *single* answer (*ans*) for every request (*req*), $s_{\mathbf{b}}$ occasionally produces *multiple* answers for a given request (see the underlined branch in the description of $s_{\mathbf{b}}$ below). Both systems terminate with *cls*.

$$\begin{aligned} s_{\mathbf{g}} &= \text{rec } x.\text{req}.\text{(ans.log.x + cls.nil)} \\ s_{\mathbf{b}} &= \text{req.rec } x.\text{(ans.(\underline{ans.x} + \text{log.req.x}) + \text{cls.nil})} \end{aligned}$$

We can specify that a request followed by two consecutive answers indicates invalid behaviour via the sHML formula φ_0 .

$$\begin{aligned} \varphi_0 &\stackrel{\text{def}}{=} \square [\text{ans}][\text{ans}]\text{ff} \\ &\stackrel{\text{def}}{=} \max X. [\text{ans}][\text{ans}]\text{ff} \wedge \bigwedge_{\alpha \in \text{ACT}} [\alpha]X \end{aligned}$$

where $\text{ACT} \stackrel{\text{def}}{=} \{\text{ans}, \text{req}, \text{cls}\}$. The above formula defines an invariant property requiring that, at every reachable state, whenever the system produces an answer, it cannot produce a subsequent answer i.e., $[\text{ans}]\text{ff}$. Using the semantics in Figure 1, one can check that $s_{\mathbf{g}} \in \llbracket \varphi_0 \rrbracket$, whereas $s_{\mathbf{b}} \notin \llbracket \varphi_0 \rrbracket$ since it exhibits the violating trace $s_{\mathbf{b}} \xrightarrow{\text{req}} \cdot \xrightarrow{\text{ans}} \cdot \xrightarrow{\text{ans}} \dots$, amongst others. \square

3 Controlled System Synthesis and Suppression Enforcement

We present the simplified models for suppression enforcement and controlled system synthesis adapted from [6] and [25], respectively. Both models describe the composite behaviour attained by the respective techniques. In suppression enforcement, the composite behaviour describes the *observable behaviour* obtained when the monitor and the SuS interact *at runtime*, while in controlled system synthesis, it describes the *structure* of the resulting controlled system obtained *statically prior to deployment*.

To enable our comparison between both approaches, we standardise the logics used in both works and restrict ourselves to sHML_{inv} , defined in Figure 2. sHML_{inv} is a strict subset of sHML which results from the intersection of sHML, used for suppression enforcement in [6], and $\text{HML}_{\text{inv}}^{\text{reach}}$, used for controlled system synthesis in [25].

Although the work on CSS in [25] assumes that systems do not perform internal τ actions and that output labels may be associated to system states, the work on RE assumes the converse. We therefore equalise the system models by working with respect to LTSs that do not associate labels to states, and do not perform τ actions. We however assume that the resulting monitored and controlled systems may still perform τ actions.

Since we do not focus on state-based properties, the removal of state labels is not a major limitation as we are only forgoing additional state information from the SuS. Although the removal of τ actions requires the SuS to be fully observable, this does not impose significant drawbacks as the work on CSS can easily be extended to allow such actions.

Despite the fact that controlled system synthesis differentiates between system actions that can be removed (controllable) and those which cannot (uncontrollable), the work on enforcement does not. This is also not a major limitation since enforcement models can easily be adapted to make such a distinction. However, in our first attempt at a comparison, we opt to simplify the models as much as possible, and so to enable our comparison we assume that every system action is controllable and can be removed and suppressed by the respective techniques.

Finally, since we do not liberally introduce constructs that are not present in the original models of [6, 25], the simplified models are just *restricted versions* of the original ones. Hence, the results proven with respect to these simplified models should either apply to the original ones or can be extended easily to the more general setting.

3.1 A Model for Suppression Enforcement

We use a simplified version of the operational model of enforcement presented in [6], which uses the transducers $m, n \in \text{TRN}$ defined in Figure 3. Transducers define *transformation pairs*, $\{\beta, \mu\}$, which intuitively state that β actions performed by the system should be transformed into μ actions. A transformation pair thus acts as a function that takes as input a system action α and transforms it into μ whenever $\alpha = \beta$. As a shorthand, we sometimes write $\{\beta\}$ in lieu of $\{\beta, \beta\}$ to signify that actions equal to β will remain unmodified.

The transition rules in Figure 3 yield an LTS with labels of the form $\alpha \blacktriangleright \mu$. Intuitively, a transition $m \xrightarrow{\alpha \blacktriangleright \mu} n$ denotes the fact that the transducer in state m *transforms* the visible action α (produced by the system) into action μ and transitions into state n . In this sense, the transducer action $\alpha \blacktriangleright \alpha$ denotes the *identity* transformation, while $\alpha \blacktriangleright \tau$ encodes the *suppression* transformation

Syntax

$$m, n \in \text{TRN} ::= \{\alpha, \mu\}.m \quad (\text{where } \mu \in \{\alpha, \tau\}) \\ | \sum_{i \in I} m_i \quad | \text{rec } x.m \quad | x$$

Dynamics

$$\text{ESEL} \frac{m_j \xrightarrow{\alpha \blacktriangleright \mu} n_j}{\sum_{i \in I} m_i \xrightarrow{\alpha \blacktriangleright \mu} n_j} \quad j \in I$$

$$\text{EREC} \frac{m\{\text{rec } x.m/x\} \xrightarrow{\alpha \blacktriangleright \mu} n}{\text{rec } x.m \xrightarrow{\alpha \blacktriangleright \mu} n}$$

$$\text{ETRN} \frac{}{\{\alpha, \mu\}.m \xrightarrow{\alpha \blacktriangleright \mu} m}$$

Instrumentation

$$\text{ITRN} \frac{s \xrightarrow{\alpha} s' \quad m \xrightarrow{\alpha \blacktriangleright \mu} n}{m[s] \xrightarrow{\mu} n[s']} \quad \text{IDEF} \frac{s \xrightarrow{\alpha} s' \quad m \xrightarrow{\alpha} \text{id}[s']}{m[s] \xrightarrow{\alpha} \text{id}[s']}$$


where $\text{id} \stackrel{\text{def}}{=} \text{rec } x. \sum_{\beta \in \text{ACT}} \{\beta\}.x$ and $m \xrightarrow{\alpha} \stackrel{\text{def}}{=} \#m', \mu \cdot m \xrightarrow{\alpha \blacktriangleright \mu} m'$.

Fig. 3 A model for transducers.

of action α . The key transition rule is ETRN . It states that the transformation-prefix transducer $\{\alpha, \mu\}.m$ can transform action α into μ , as long as the specifying action α is the same as the action performed by the system. In this case, the transformed action is μ , and the transducer state that is reached is m .

The remaining rules ESEL and EREC respectively define the standard selection and recursion operations. A sum of transducers $\sum_{i \in I} m_i$ (where I is a finite set) can reduce via ESEL to some n_j over some action $\alpha \blacktriangleright \mu$, whenever there exists a transducer m_j in the summation that reduces to n_j over the same action. Rule EREC enables a recursion transducer $\text{rec } x.m$ to reduce to some n when its unfolded instance $m\{\text{rec } x.m/x\}$ reduces to n as well. We encode the identity monitor, id , and the suppression monitor, sup , as $\text{rec } x. \sum_{\beta \in \text{ACT}} \{\beta\}.x$ and $\text{rec } x. \sum_{\beta \in \text{ACT}} \{\beta, \tau\}.x$ respectively, i.e., as recursive monitors respectively defining an identity and suppression transformation for every possible action $\beta \in \text{ACT}$ that can be performed by the system.

Figure 3 also describes an *instrumentation* relation, which composes the behaviour of the SuS s with the transformations of a transducer monitor m that *agrees* with the (observable) actions ACT of s . The term $m[s]$ thus denotes the resulting *monitored system* whose transitions are labelled with actions in $\text{ACT} \cup \{\tau\}$ from the system's LTS. Concretely, rule ITRN states that when a system s transitions with an observable action α to s' and the transducer m can *transform* this action into μ and transition to n , the instrumented system $m[s]$ transitions with action μ to $n[s']$. Rule IDEF is analogous to standard monitor instrumentation rules for



$$\text{where } s_b \stackrel{\text{def}}{=} \text{req}.s_b^1 \quad s_b^1 \stackrel{\text{def}}{=} \text{ans}.s_b^2 + \text{cls}.nil \\ s_b^2 \stackrel{\text{def}}{=} \text{ans}.s_b^1 + \text{log}.s_b^3 \quad s_b^3 \stackrel{\text{def}}{=} \text{req}.s_b^1.$$

Fig. 4 The runtime execution graph of the monitored system.

premature termination of the transducer [5, 19, 20, 22], and accounts for underspecification of transformations. Thus, if a system s transitions with an observable action α to s' , and the transducer m does not specify how to transform that action ($m \not\xrightarrow{\alpha}$), the system is still allowed to transition while the transducer defaults to acting like the identity monitor, id , from that point onwards.

Example 2 Consider the suppression transducer m_s below:

$$m_s \stackrel{\text{def}}{=} \text{rec } x. (\{\text{ans}\}.m'_s + \{\text{req}\}.x + \{\text{cls}\}.x + \{\text{log}\}.x) \\ m'_s \stackrel{\text{def}}{=} \{\text{ans}, \tau\}.\text{sup} + \{\text{req}\}.x + \{\text{cls}\}.x + \{\text{log}\}.x$$

where sup recursively suppresses every action $\beta \in \text{ACT}$ that can be performed by the system from that point onwards. When instrumented with system s_b from Example 1, the monitor prevents the monitored system $m_s[s_b]$ from answering twice in a row by suppressing the second answer and every subsequent visible action:

$$m_s[s_b] \xrightarrow{\text{req.ans}} \cdot \xrightarrow{\tau} \text{sup}[s_b].$$

When equipped with this dynamic action suppression mechanism, the resulting monitored system $m_s[s_b]$ never violates formula φ_0 at runtime – this is illustrated by the runtime execution graph in Figure 4. \square

3.2 Synthesising Controlled Systems

Figure 5 presents a synthesis function that takes a system $\langle \text{SYS}, \text{ACT}, \rightarrow \rangle$ and a formula $\varphi \in \text{SHML}$ and constructs a controlled version of the system that satisfies the formula. The new system is synthesised in two stages. In the first stage, a new transition relation $\mapsto_{\subseteq} (\text{SYS} \times \text{SHML}) \times \text{ACT} \times (\text{SYS} \times \text{SHML})$ is constructed over the state-formula product space, $(\text{SYS} \times \text{SHML})$. Intuitively, a state (s, φ) represents a version of state s that is controlled according to the constraints required by φ . A μ -transition $(s, \varphi) \xrightarrow{\mu} (s', \varphi')$ represents a reduction from one controlled state to another. The transition relation associates a SHML formula to the initial system state and defines how this changes when the system transitions to other subsequent states. The

Static Composition

$$\begin{array}{c}
\text{cBOOL} \frac{s \xrightarrow{\alpha} s' \quad b \in \{\text{tt}, \text{ff}\}}{(s, b) \xrightarrow{\alpha} (s', b)} \\
\text{cNEC1} \frac{s \xrightarrow{\alpha} s'}{(s, [\alpha]\varphi) \xrightarrow{\alpha} (s', \varphi)} \\
\text{cNEC2} \frac{s \xrightarrow{\beta} s' \quad \beta \neq \alpha}{(s, [\alpha]\varphi) \xrightarrow{\beta} (s', \text{tt})} \\
\text{cAND} \frac{(s, \varphi) \xrightarrow{\alpha} (s', \varphi') \quad (s, \psi) \xrightarrow{\alpha} (s', \psi')}{(s, \varphi \wedge \psi) \xrightarrow{\alpha} (s', \min(\varphi' \wedge \psi'))} \\
\text{cMAX} \frac{(s, \varphi \{\max X. \varphi / X\}) \xrightarrow{\alpha} (s', \psi)}{(s, \max X. \varphi) \xrightarrow{\alpha} (s', \min(\psi))}
\end{array}$$

Synthesizability Test

$$\frac{\psi \in \{\text{tt}, X, [\alpha]\varphi\} \quad (s, \varphi) \downarrow \psi}{(s, \varphi) \downarrow \psi} \quad \frac{(s, \varphi) \downarrow \psi_1 \quad (s, \varphi) \downarrow \psi_2}{(s, \varphi) \downarrow (\psi_1 \wedge \psi_2)}$$

$$\frac{(s, \varphi) \downarrow \psi}{(s, \varphi) \downarrow \max X. \psi}$$

Invalid Transition Removal

$$\text{cTR} \frac{(s, \varphi) \xrightarrow{\alpha} (s', \varphi') \quad (s', \varphi') \downarrow \varphi'}{(s, \varphi) \xrightarrow{\alpha} (s', \varphi')}$$

Fig. 5 The Controlled System Synthesis.

composite behaviour of the formula and the system is statically computed using the first five rules in Figure 5.

cBOOL adds a transition from (s, b) when the formula is $b \in \{\text{tt}, \text{ff}\}$ if that transition is possible in s . Rules cNEC1 and cNEC2 add a transition from $[\alpha]\varphi$ to φ when s has a transition over α , and to tt if it reduces over $\beta \neq \alpha$. cAND adds a transition for conjunct formulas, $\varphi \wedge \psi$, when both formulas can reduce independently to some φ' and ψ' , with the formula of the end state of the new transition being $\min(\varphi' \wedge \psi')$. Finally, cMAX adds a fixpoint $\max X. \varphi$ transition to $\min(\psi)$, when its unfolding can reduce to ψ . In both cAND and cMAX, $\min(\varphi)$ stands for a *minimal* logically equivalent formula of φ . This is an oversimplification of the syntactic manipulation techniques used in [25] to avoid synthesising an infinite LTS due to invariant formulas and conjunctions, see [25] for more details.

Example 3 Formulas $\varphi' \wedge \text{tt}$, $\varphi' \wedge \text{ff}$ and $\varphi \wedge \psi \wedge \psi$ are *logically equivalent* to (and can thus be minimized into) φ' , ff and $\varphi \wedge \psi$ respectively. \square

Instead of defining a rule for fixpoints, the authors of [25] define a synthesis rule directly for invariance stating that when $(s, \varphi) \xrightarrow{\alpha} (s', \varphi')$, then $(s, \square \varphi) \xrightarrow{\alpha}$

$(s', \min(\square \varphi \wedge \varphi'))$. We, however, opted to generalize this rule to fixpoints to simplify our comparison, while still limiting ourselves to sHML_{inv} formulas. This is possible since by encoding $\square \varphi$ as $\max X. \varphi \wedge \bigwedge_{\beta \in \text{ACT}} [\beta] X$, we get that $(s, \max X. \varphi \wedge \bigwedge_{\beta \in \text{ACT}} [\beta] X) \xrightarrow{\alpha} (s', \min(\psi))$ when $(s, \varphi) \xrightarrow{\alpha} (s', \varphi')$ where $\psi \stackrel{\text{def}}{=} (\max X. \varphi \wedge \bigwedge_{\beta \in \text{ACT}} [\beta] X) \wedge \varphi'$ and $\min(\psi)$ is the encoded version of $\min(\square \varphi \wedge \varphi')$.

The second stage of the synthesis involves using rule cTR to remove invalid transitions that lead to violating states; this yields the required transition function for the controlled system. This rule relies on the synthesizability test rules to tell whether a controlled state, (s, φ) , is valid or not. Intuitively, the test rules fail whenever the current formula φ is semantically equivalent to ff , e.g., formulas $\max X. ([\alpha] X \wedge \text{ff})$ and $\varphi \wedge \text{ff}$ both fail the synthesizability test rules as they are equivalent to ff . Concretely, the test is vacuously satisfied by truth, tt , logical variables, X , and guarded formulas, $[\alpha]\varphi$, as none of them are logically equivalent to ff . Conjunct formulas, $\psi_1 \wedge \psi_2$, pass the test when both ψ_1 and ψ_2 pass independently. A fixpoint, $\max X. \varphi'$, is synthesisable if φ' passes the test.

Transitions that lead to a state that fails the test are therefore removed, and transitions outgoing from failing states become redundant as they are unreachable. The resulting transition function is then used to construct the controlled LTS $\langle (\text{Sys} \times \text{sHML}_{\text{inv}}), \text{ACT}, \rightarrow \rangle$.

Example 4 From φ_0 and $s_{\mathbf{b}}$ of Example 1 we can synthesise a controlled system in two stages. In the first stage we compose them together using the composition rules of Figure 5. We start by generating the composite transition $(s_{\mathbf{b}}, \varphi_0) \xrightarrow{\text{req}} (s_{\mathbf{b}}^1, \varphi_0)$ via rules cMAX and cNEC since $s_{\mathbf{b}} \xrightarrow{\text{req}} s_{\mathbf{b}}^1$, and keep on applying the respective rules to the rest of $s_{\mathbf{b}}$'s transitions until we obtain the LTS of Figure 6. The (grey) ans transition leading to the test failing state, $(s_{\mathbf{b}}, \text{ff}) \downarrow$, is then removed in the second stage along with its outgoing (grey) transitions, therefore generating the required (black) controlled system. \square

4 Establishing a static counterpart to enforcement

To be able to establish whether CSS is a static counterpart to suppression enforcement, we must first formalise the meaning of a “static counterpart”. We thus define it as Definition 3.

Definition 3 (Static Counterpart) A static verification technique \mathcal{S} is the *static counterpart* of suppression enforcement (in the context of safety properties) when, for every LTS $(\text{Sys}, \text{ACT}, \rightarrow)$, formula $\varphi \in \text{sHML}_{\text{inv}}$ and




Fig. 6 The LTS obtained from controlling s_b via φ_5 .

$s \in \text{SYS}$, there exists a transducer m so that $m[s] \in \llbracket \varphi \rrbracket$ iff $\mathcal{S}(s) \in \llbracket \varphi \rrbracket$ (where $\mathcal{S}(s)$ is a statically reformulated version of s obtained from applying \mathcal{S}). \square

Determining that CSS is the static counterpart of suppression enforcement (as stated by Definition 3) is inherently difficult as it requires showing that every SHML_{inv} formula (of which there is an infinite amount) can be enforced using both techniques. Despite this, Examples 2 and 4 already provide the intuition that there exists some level of correspondence between these two techniques. In fact, from the monitored execution graph of Figure 4 and the controlled LTS in Figure 6 one can notice that they both execute the *same set of traces*, and are therefore *trace equivalent*. Since trace equivalent systems satisfy the same set of safety properties (Theorem 1), establishing trace equivalence suffices to conclude that the controlled LTS is statically achieving the same result obtained dynamically by the monitored one, and that it is therefore its static counterpart.

We thus begin by showing that trace equivalent systems satisfy the same set of *safety properties*. As the (recursion-free) subset of SHML characterises regular safety properties [22], this means that systems sharing the same traces also satisfy the same SHML formulas.

Theorem 1 *Let s and r be system states in an LTS. Then $\text{traces}(s) = \text{traces}(r)$ iff s and r satisfy exactly the same SHML formulas.* \square

The proof for this theorem we present relies on the work on detection (runtime verification) monitors by Francalanza *et al.* in [22]. Detection monitors d in [22] can reject a trace t at runtime by issuing the verdict **no** whenever they detect that an SHML formula has been violated by t , i.e., $d \xrightarrow{t} \text{no}$. In respect to these detection monitors, Francalanza *et al.* prove the following results.

- **Detection Soundness:** For every SHML formula φ , system s , trace t and detection monitor d , if $s \xrightarrow{t}$ and $d \xrightarrow{t} \text{no}$ then $s \notin \llbracket \varphi \rrbracket$.
- **Detection Completeness:** For every SHML formula φ , system s , if $s \notin \llbracket \varphi \rrbracket$ then there exists a trace t and a detection monitor d such that $s \xrightarrow{t}$ and $d \xrightarrow{t} \text{no}$.

Detection soundness states that if a system state s executes a trace t that gets rejected by a detection

monitor d , then s violates φ . Completeness states the converse. Using this framework we can now easily prove Theorem 1 as follows.

Proof for Theorem 1. Assume that

$$\text{traces}(s) \subseteq \text{traces}(r) \quad \text{and that} \quad (1)$$

$$s \notin \llbracket \varphi \rrbracket. \quad (2)$$

Knowing (2) and *Detection Completeness* from [22], we can infer that there exists a trace t , and detection monitor d , such that when s executes t , d rejects it for being invalid, i.e., $d \xrightarrow{t} \text{no}$. Hence, since from (1) we know that the invalid trace t can also be executed by r , by *Detection Soundness* from [22] we can also conclude that $r \notin \llbracket \varphi \rrbracket$ as required, and we are done. \square

Hence, since trace equivalent systems satisfy the same set of safety properties (Theorem 1), it suffices to conclude that the controlled LTS can produce the same set of traces as that generated by a monitored one at runtime.

Theorem 2 (Trace Equivalence) *For every formula $\varphi \in \text{SHML}_{\text{inv}}$, there exists a monitor m such that for every $s \in \text{SYS}$, $\text{traces}(m[s]) = \text{traces}(s, \varphi)$.* \square

The existential quantification on the monitor m in Theorem 2 entails the need of using some sort of mapping that maps SHML_{inv} formulas to suppression monitors. To be able to prove this result, we thus define a function that maps SHML_{inv} formulas to enforcement transducers. We reduce the complexity of this mapping by defining it over the normalised SHML formulas instead.

Definition 4 (sHML normal form) The set of normalised SHML formulas is defined as:

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

In addition, a normalised SHML formula φ must satisfy the following conditions:

1. In each subformula of φ of the form $\bigwedge_{i \in I} [\alpha_i] \varphi_i$, the α_i 's are *pairwise different*, denoted as $\#_{i \in I} \alpha_i$, i.e., $\forall i, j \in I \cdot \text{if } i \neq j \text{ then } \alpha_i \neq \alpha_j$.
2. For every $\max X. \varphi$ we have $X \in \text{fv}(\varphi)$.
3. Every logical variable is *guarded* by a modal necessity. \square

In previous work [4, 6] we proved that, despite being a syntactic subset of SHML, SHML_{nf} is *semantically equivalent* to SHML. Hence, since SHML_{inv} is a (strict) subset of SHML, for every SHML_{inv} formula we can always find an equivalent SHML_{nf} formula. This means that by defining our mapping function in terms of SHML_{nf}, we can still map every formula in SHML_{inv} to the respective monitor.

We proceed to define our mapping function over normalised SHML formulas.

Definition 5 We inductively define the mapping function $\llbracket - \rrbracket : \text{SHML}_{\text{nf}} \mapsto \text{TRN}$ as follows:

$$\begin{aligned} \llbracket X \rrbracket &\stackrel{\text{def}}{=} x & \llbracket \text{tt} \rrbracket &\stackrel{\text{def}}{=} \text{id} & \llbracket \text{ff} \rrbracket &\stackrel{\text{def}}{=} \text{sup} \\ \llbracket \max X.\varphi \rrbracket &\stackrel{\text{def}}{=} \text{rec } x.\llbracket \varphi \rrbracket & \llbracket \bigwedge_{i \in I} [\alpha_i]\varphi_i \rrbracket &\stackrel{\text{def}}{=} \sum_{i \in I} m_i \end{aligned}$$

$$\text{where } m_i \stackrel{\text{def}}{=} \begin{cases} \{\alpha_i, \alpha_i\}.\llbracket \varphi_i \rrbracket & \text{if } \varphi_i \neq \text{ff} \\ \{\alpha_i, \tau\}.\llbracket \text{ff} \rrbracket & \text{otherwise} \end{cases} \quad \square$$

The function is compositional. It assumes a bijective mapping between fixpoint variables and monitor recursion variables and converts logical variables X accordingly, whereas maximal fixpoints, $\max X.\varphi$, are converted into the corresponding recursive monitor. The function also converts truth and falsehood formulas, tt and ff , into the identity monitor id and the suppression monitor sup respectively. Normalized conjunctions, $\bigwedge_{i \in I} [\alpha_i]\varphi_i$, are mapped into a *summation* of monitors, $\sum_{i \in I} m_i$, where every branch m_i can be either prefixed by an identity transformation when $\varphi_i \neq \text{ff}$, or by a suppression transformation otherwise.

Notice that the requirement that $\varphi_i \neq \text{ff}$ in Definition 5 is in some sense analogous to the pruning of transitions applied by the CSS rule cTR of Figure 5 to retain the valid transitions only. In this mapping function, this requirement is essential to ensure that only the actions that do not lead to violations of the input formula remain unsuppressed by the resulting monitor.

Example 5 Recall formula φ_0 from Example 1 which can be normalised as:

$$\begin{aligned} \varphi_0 &\stackrel{\text{def}}{=} \max X. [\text{ans}]\varphi'_0 \wedge [\text{req}]X \wedge [\text{log}]X \wedge [\text{cls}]X \\ \varphi'_0 &\stackrel{\text{def}}{=} [\text{ans}]\text{ff} \wedge [\text{req}]X \wedge [\text{log}]X \wedge [\text{cls}]X. \end{aligned}$$

Using the mapping function defined in Definition 5, we generate monitor

$$\begin{aligned} \llbracket \varphi_0 \rrbracket &= \text{rec } x.\{\text{ans}\}.\llbracket \varphi'_0 \rrbracket + \{\text{req}\}.x + \{\text{log}\}.x + \{\text{cls}\}.x \\ \llbracket \varphi'_0 \rrbracket &= \{\text{ans}, \tau\}.\text{sup} + \{\text{req}\}.x + \{\text{log}\}.x + \{\text{cls}\}.x \end{aligned}$$

which is identical to m_s from Example 2. \square

With this mapping function in hand, we are able to prove Theorem 2 as a corollary of Proposition 1.

Proposition 1 For every LTS $\langle \text{SYS}, \text{ACT}, \rightarrow \rangle$, SHML_{nf} formula φ , $s \in \text{SYS}$ and trace t , it holds that $t \in \text{traces}(m[s])$ iff $t \in \text{traces}((s, \varphi))$ where $\llbracket \varphi \rrbracket = m$. \square

In our proof we rely on Lemma 1.

Lemma 1 For every system s and r , SHML_{nf} formula φ and action α , if $\llbracket \varphi \rrbracket[s] \xrightarrow{\alpha} r$ then $\llbracket \varphi \rrbracket[s] \xrightarrow{\alpha} r$. \square

The proof for this lemma is provided after that of Proposition 1. We now proceed to prove the *if* and *only-if* cases separately as follows.

Proof for the only-if case. We proceed by induction on the length of t .

Case $t = \varepsilon$. This case holds vacuously since the empty trace can be executed by every system, that is, $\varepsilon \in \text{traces}((s, \varphi))$ as required.

Case $t = \alpha t'$. Assume that $\alpha t' \in \text{traces}(\llbracket \varphi \rrbracket[s])$ and so by the definition of *traces* we know that there exists a system r such that

$$t' \in \text{traces}(r) \quad (3)$$

and that $\llbracket \varphi \rrbracket[s] \xrightarrow{\alpha} r$. Hence, by Lemma 1 we get that

$$\llbracket \varphi \rrbracket[s] \xrightarrow{\alpha} r. \quad (4)$$

We now proceed by case analysis on φ .

- $\varphi \in \{X, \text{ff}\}$: These cases do not apply since they contradict assumption (4), namely since $\nexists m \cdot \llbracket X \rrbracket = m$, and since $\forall t \in \text{ACT} \cdot \llbracket \text{ff} \rrbracket[s] \not\xrightarrow{t}$ where $\llbracket \text{ff} \rrbracket = \text{sup}$.
- $\varphi = \text{tt}$: Since $\llbracket \text{tt} \rrbracket = \text{id}$, by rule ITRN , from (4) we get that

$$s \xrightarrow{\alpha} s' \quad (5)$$

and that $r = \text{id}[s'] = \llbracket \text{tt} \rrbracket[s']$ which in conjunction with (3) and the *inductive hypothesis* we can deduce that

$$t' \in \text{traces}((s', \text{tt})). \quad (6)$$

Since $\varphi = \text{tt}$ and knowing (5) we can synthesise the controlled transition $(s, \text{tt}) \xrightarrow{\alpha} (s', \text{tt})$ so that from (6) we conclude that $\alpha t' \in \text{traces}((s, \text{tt}))$ as required.

- $\varphi = \bigwedge_{i \in I} [\alpha_i]\varphi_i$: Since $\llbracket \bigwedge_{i \in I} [\alpha_i]\varphi_i \rrbracket$ synthesises the monitor $\left(\sum_{i \in I} \begin{cases} \{\alpha_i, \alpha_i\}.\llbracket \varphi_i \rrbracket & \text{if } \varphi_i \neq \text{ff} \\ \{\alpha_i, \tau\}.\llbracket \text{ff} \rrbracket & \text{otherwise} \end{cases} \right)$, we must explore the instrumentation rules that permit for the α -reduction in (4).
- ITRN : By applying rule ITRN to (4) we have that

$$s \xrightarrow{\alpha} s' \quad (7)$$

$$r = m[s'] \quad (8)$$

and that $\left(\sum_{i \in I} \begin{cases} \{\alpha_i, \alpha_i\} \cdot \langle \varphi_i \rangle & \text{if } \varphi_i \neq \text{ff} \\ \{\alpha_i, \tau\} \cdot \langle \text{ff} \rangle & \text{otherwise} \end{cases} \right) \xrightarrow{\alpha \blacktriangleright \alpha}$

m so that by rules ESEL and ETRN we know that

$$\exists j \in I \cdot \alpha_j = \alpha \quad (9)$$

$$m = \langle \varphi_j \rangle \text{ (where } \varphi_j \neq \text{ff}). \quad (10)$$

Knowing (7), (9) and that $\varphi_j \neq \text{ff}$ we can synthesise the controlled transition

$$(s, [\alpha_j] \varphi_j) \xrightarrow{\alpha} (s', \varphi_j). \quad (11)$$

Moreover, since the conjunct modal necessities are *pairwise disjoint*, from (9) we infer that for every $i \in I \setminus \{j\}$, $\alpha_i \neq \alpha$ and so we can synthesise the controlled transition $(s, [\alpha_i] \varphi_i) \xrightarrow{\alpha} (s', \text{tt})$ which in conjunction with (11) can be synthesised as the transition

$$(s, \bigwedge_{i \in I} [\alpha_i] \varphi_i) \xrightarrow{\alpha} (s', \varphi_j) \quad (12)$$

since $\min(\varphi_j \wedge \bigwedge_{i \in I \setminus \{j\}} \text{tt}) = \varphi_j$. Finally, since by (3), (8), (10) and the *inductive hypothesis* we have that $t' \in \text{traces}((s', \varphi_j))$, from (12) we conclude that $\alpha t' \in \text{traces}((s, \bigwedge_{i \in I} [\alpha_i] \varphi_i))$ as required.

– IDEF: By rule IDEF we get that

$$s \xrightarrow{\alpha} s' \quad (13)$$

$$r = \text{id}[s'] \quad (14)$$

and that $\left(\sum_{i \in I} \begin{cases} \{\alpha_i, \alpha_i\} \cdot \langle \varphi_i \rangle & \text{if } \varphi_i \neq \text{ff} \\ \{\alpha_i, \tau\} \cdot \langle \text{ff} \rangle & \text{otherwise} \end{cases} \right) \xrightarrow{\alpha}$ from which we can infer that for every $i \in I$, $\alpha_i \neq \alpha$. With this result, from (13) we can thus synthesise the controlled transition $(s, \bigwedge_{i \in I} [\alpha_i] \varphi_i) \xrightarrow{\alpha} (s', \text{tt})$ and so since $\langle \text{tt} \rangle = \text{id}$ and by (3), (14) and the *inductive hypothesis* we have that $t' \in \text{traces}((s', \text{tt}))$. Hence, we can conclude that $\alpha t' \in \text{traces}((s, \bigwedge_{i \in I} [\alpha_i] \varphi_i))$ as required.

– $\varphi = \max X. \varphi'$: Since $X \in \mathbf{fv}(\varphi')$ we can deduce that $\varphi' \notin \{\text{tt}, \text{ff}\}$, and also that $\varphi' \neq X$ since logical variables are required to be guarded in SHML_{nf}. We can thus infer that φ' adheres to a specific structure, that is, $\max Y_0 \dots Y_n. \bigwedge_{i \in I} [\alpha_i] \varphi_i$ (where $\max Y_0 \dots Y_n$ is an arbitrary number of fixpoint declarations, possibly none). Hence, since from (4) we can also infer that $\langle \max X. \max Y_0 \dots Y_n. \bigwedge_{i \in I} [\alpha_i] \varphi_i \rangle [s] \xrightarrow{\alpha} r$, and since fixpoint unfolding preserves semantics, we get that

$$\langle \bigwedge_{i \in I} [\alpha_i] \varphi_i \{ \dots \} \rangle [s] \xrightarrow{\alpha} r$$

where $\{ \dots \} \stackrel{\text{def}}{=} \{ \max XY_0 \dots Y_n. \bigwedge_{i \in I} [\alpha_i] \varphi_i / X, \dots \}$ (15)

After reaching the point where we know (15), the proof proceeds as per the previous case (*i.e.*, when $\varphi = \bigwedge_{i \in I} [\alpha_i] \varphi_i$). We thus skip this part of the proof and

simply deduce that $\alpha t' \in \text{traces}((s, \bigwedge_{i \in I} [\alpha_i] \varphi_i \{ \dots \}))$. Hence, since unfolded recursive formulas are equivalent to their folded versions, and since φ' is defined as $\max Y_0 \dots Y_n. \bigwedge_{i \in I} [\alpha_i] \varphi_i$, we can thus deduce that $\alpha t' \in \text{traces}((s, \max X. \varphi'))$ as required, and so we are done. \square

Proof for the if case. We again proceed by induction on the structure of t .

Case $t = \varepsilon$. This case holds vacuously since $\varepsilon \in \text{traces}(\langle \varphi \rangle [s])$ as required.

Case $t = \alpha t'$. Assume that $\alpha t' \in \text{traces}((s, \varphi))$ and so by the definition of *traces* we know that there exists a system r such that

$$(s, \varphi) \xrightarrow{\alpha} r \quad (16)$$

$$t' \in \text{traces}(r). \quad (17)$$

We proceed by case analysis on φ .

- $\varphi \in \{\text{ff}, X\}$: These cases do not apply because state (s, φ) is invalid.
- $\varphi = \text{tt}$: Since $(s, \text{tt}) \xrightarrow{\alpha} (s', \text{tt})$ this case holds trivially since $r = (s', \text{tt})$ and so by (17) and the *inductive hypothesis* we get that $t' \in \text{traces}(\langle \text{tt} \rangle [s'])$ and since $\langle \text{tt} \rangle = \text{id}$, by rules ITRN and ETRN we have that $\langle \text{tt} \rangle [s] \xrightarrow{\alpha} \langle \text{tt} \rangle [s']$ which allows us to conclude that $\alpha t' \in \text{traces}(\langle \text{tt} \rangle [s])$.
- $\varphi = \bigwedge_{i \in I} [\alpha_i] \varphi_i$ and $\not\#_{i \in I} \alpha_i$: In this case we have that

$$(s, \bigwedge_{i \in I} [\alpha_i] \varphi_i) \xrightarrow{\alpha} r \quad (18)$$

$$\exists \psi \cdot r = (s', \min(\psi)) \quad (19)$$

and so since the branches of the conjunction are disjoint, we only need to further investigate the following cases:

- $\forall i \in I \cdot \alpha_i \neq \alpha$: Hence, from (18) and by rule CNEC2 we can infer that for every $i \in I$ we have that $(s, [\alpha_i] \varphi_i) \xrightarrow{\alpha} (s', \text{tt})$ and that

$$s \xrightarrow{\alpha} s' \quad (20)$$

$$\min(\psi) = \text{tt} \text{ (since } \psi = \bigwedge_{i \in I} \text{tt}). \quad (21)$$

Therefore, as we know (20) and that for every $i \in I$ then $\alpha_i \neq \alpha$, by rules ETRN and ESEL we can infer that $\left(\sum_{i \in I} \begin{cases} \{\alpha_i, \alpha_i\} \cdot \langle \varphi_i \rangle & \text{if } \varphi_i \neq \text{ff} \\ \{\alpha_i, \tau\} \cdot \langle \text{ff} \rangle & \text{otherwise} \end{cases} \right) \xrightarrow{\alpha}$ and so by the definition of $\langle - \rangle$ and rule IDEF we conclude that

$$\langle \bigwedge_{i \in I} [\alpha_i] \varphi_i \rangle [s] \xrightarrow{\alpha} \langle \text{tt} \rangle [s']. \quad (22)$$

Finally, from (17), (19), (21) and by the *inductive hypothesis* we have that $t' \in \text{traces}(\llbracket \text{tt} \rrbracket [s'])$ and so from (22), we infer that $\alpha t' \in \text{traces}(\llbracket \bigwedge_{i \in I} [\alpha_i] \varphi_i \rrbracket [s])$.

- $\exists j \in I \cdot \eta_j = \alpha$ but $\forall i \in I \setminus \{j\} \cdot \alpha_i \neq \alpha$: In this case, from the controlled synthesis rules and from (18) and (19) we can infer that $\exists j \in I \cdot (s, [\alpha_j] \varphi_j) \xrightarrow{\alpha} (s', \varphi_j)$ and that $\forall i \in I \setminus \{j\} \cdot (s, [\alpha_i] \varphi_i) \xrightarrow{\alpha} (s', \text{tt})$ and finally that

$$s \xrightarrow{\alpha} s' \quad (23)$$

$$\min(\psi) = \min(\varphi_j \wedge \bigwedge_{i \in I} \text{tt}) = \varphi_j \quad (24)$$

where $\varphi_j \neq \text{ff}$, as otherwise, if $\varphi_j = \text{ff}$ the resulting state $(s', \min(\text{ff}))$ would be invalid and thus removed by the synthesis along with any transitions leading to it (including (18)). Knowing that there exists $j \in I$ so that $\alpha_j \neq \alpha$ and by rule ETRN we can also deduce that $\{\alpha_j, \alpha_j\} \cdot (\varphi_j) \xrightarrow{\alpha} (\varphi_j)$ and so by rule ESEL we have that $\left(\sum_{i \in I} \begin{cases} \{\alpha_i, \alpha_i\} \cdot (\varphi_i) & \text{if } \varphi_i \neq \text{ff} \\ \{\alpha_i, \tau\} \cdot (\text{ff}) & \text{otherwise} \end{cases} \right) \xrightarrow{\alpha} (\varphi_j)$. This means that by (23), rule ITRN and the definition of $\llbracket - \rrbracket$ we conclude that

$$\llbracket \bigwedge_{i \in I} [\alpha_i] \varphi_i \rrbracket [s] \xrightarrow{\alpha} \llbracket \varphi_j \rrbracket [s']. \quad (25)$$

Finally, since from (17), (19), (24) and the *inductive hypothesis* we know that $t' \in \text{traces}(\llbracket \varphi_j \rrbracket [s'])$, from (25) we can infer that $\alpha t' \in \text{traces}(\llbracket \bigwedge_{i \in I} [\alpha_i] \varphi_i \rrbracket [s])$ as required.

- $\varphi = \max X.\varphi'$ and $X \in \mathbf{fv}(\varphi')$: We now have that $(s, \max X.\varphi') \xrightarrow{\alpha} (s', \psi)$ because

$$(s, \varphi' \{ \max X.\varphi' / X \}) \xrightarrow{\alpha} (s', \psi) \quad (26)$$

and so since φ' can neither be X (since SHML_{nf} requires fixpoint variables to be guarded) nor ff or tt (since $X \in \mathbf{fv}(\varphi')$) we can deduce that φ' must have the form $\max Y_0 \dots Y_n \cdot \bigwedge_{i \in I} [\alpha_i] \varphi_i$ and so since fixpoint unfolding preserves formula semantics, from (26) we can subsequently deduce that $(s, \bigwedge_{i \in I} [\alpha_i] \varphi_i \{ \dots \}) \xrightarrow{\alpha} (s', \psi)$ where $\{ \dots \} \stackrel{\text{def}}{=} \{ \max X Y_0 \dots Y_n \cdot \bigwedge_{i \in I} [\alpha_i] \varphi_i / X, \dots \}$. From this point onwards the proof proceeds as per the previous case ($\varphi = \bigwedge_{i \in I} [\alpha_i] \varphi_i$), we thus skip this part of the proof and safely conclude that

$$\alpha t' \in \text{traces}(\llbracket \bigwedge_{i \in I} [\alpha_i] \varphi_i \{ \dots \} \rrbracket [s]). \quad (27)$$

Since fixpoint folding preserves semantics and $\varphi' = \max Y_0 \dots Y_n \cdot \bigwedge_{i \in I} [\alpha_i] \varphi_i$, from (27) we thus conclude that $\alpha t' \in \text{traces}(\llbracket \max X.\varphi' \rrbracket [s])$ as required, and so we are done. \square

We now prove the auxiliary lemma Lemma 1 as follows. The reader may safely skip the following proofs upon first reading and proceed from Page 12.

Proof for Lemma 1. We must prove that for every system state s and r , SHML_{nf} formula φ and action α , if $\llbracket \varphi \rrbracket [s] \xRightarrow{\alpha} r$ then $\llbracket \varphi \rrbracket [s] \xrightarrow{\alpha} r$.

Since we assume that the SuS s does not perform τ actions, by the rules in our enforcement model we know that the only case when a τ reduction is part of a monitored execution occurs when the monitor suppresses a (visible) action of s . We proceed by case analysis on φ .

Case $\varphi \in \{X, \text{ff}\}$. These cases do not apply since $\#m \cdot \llbracket X \rrbracket = m$ and since $\llbracket \text{ff} \rrbracket = \text{sup}$ and so $\# \beta \in \text{ACT} \cdot \text{sup}[s] \xrightarrow{\beta}$.

Case $\varphi = \text{tt}$. Since $\llbracket \text{tt} \rrbracket = \text{id}$ cannot suppress any action, we deduce that the weak transition in $(\llbracket \text{tt} \rrbracket, s) \xRightarrow{\alpha} r$ is in fact a strong one and so that $(\llbracket \text{tt} \rrbracket, s) \xrightarrow{\alpha} r$ as required.

Case $\varphi = \bigwedge_{i \in I} [\alpha_i] \varphi_i$. Assume that

$$\left(\sum_{i \in I} \begin{cases} \{\alpha_i\} \cdot (\varphi_i) & \text{if } \varphi_i \neq \text{ff} \\ \{\alpha_i, \tau\} \cdot (\text{ff}) & \text{otherwise} \end{cases} \right) [s] \xrightarrow{\alpha} r. \quad (28)$$

From the weak reduction in (28) we infer that the system must perform some action β which is then suppressed by one of the monitor's branches, and so there must exist an index $j \in I$ so that $\alpha_j = \beta$ and $\{\alpha_j, \tau\} \cdot (\text{ff}) \xrightarrow{\beta, \tau} (\text{ff})$. However, since $\llbracket \text{ff} \rrbracket = \text{sup}$, we know that if any invalid action β were to be executed by s and, as a consequence, suppressed by the monitor, any subsequent action (including α) would also be suppressed by sup , in which case the instrumented system in (28) would be unable to eventually execute α and thus yield a contradiction. Therefore, the only way that transition (28) can happen is when the monitor *does not* suppress any action prior to executing α , which thus means that the weak reduction in (28) is in fact a strong one, i.e.,

$$\left(\sum_{i \in I} \begin{cases} \{\alpha_i\} \cdot (\varphi_i) & \text{if } \varphi_i \neq \text{ff} \\ \{\alpha_i, \tau\} \cdot (\text{ff}) & \text{otherwise} \end{cases} \right) [s] \xrightarrow{\alpha} r \text{ as required.}$$

Case $\varphi = \max X.\varphi'$ where $X \in \mathbf{fv}(\varphi')$. Assume that $\llbracket \max X.\varphi' \rrbracket [s] \xRightarrow{\alpha} r$ and so since $\llbracket \max X.\varphi' \rrbracket$ is logically equivalent to $\llbracket \varphi' \{ \max X.\varphi' / X \} \rrbracket$ we can deduce that

$$\llbracket \varphi' \{ \max X.\varphi' / X \} \rrbracket [s] \xRightarrow{\alpha} r. \quad (29)$$

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 form of

$$\varphi' = \max Y_0 \dots Y_n \cdot \bigwedge_{i \in I} [\alpha_i] \varphi_i \quad (30)$$

where $\max Y_0 \dots Y_n \dots$ represents an arbitrary number of fixpoint declarations, possibly none. Hence, since $\llbracket \varphi' \rrbracket$ is logically equivalent to $\llbracket \bigwedge_{i \in I} [\alpha_i] \varphi_i \{..\} \rrbracket$ where $\{..\} \stackrel{\text{def}}{=} \{\max X Y_0 \dots Y_n. \bigwedge_{i \in I} [\alpha_i] \varphi_i / X, \dots\}$, from (29) and (30) we have that

$$\llbracket \bigwedge_{i \in I} [\alpha_i] \varphi_i \{..\} \rrbracket [s] \xrightarrow{\alpha} r. \quad (31)$$

Having reached the point where we know (31), the proof becomes identical as per the previous case ($\varphi = \bigwedge_{i \in I} [\alpha_i] \varphi_i$), we thus skip this part of the proof and safely conclude that $\llbracket \bigwedge_{i \in I} [\alpha_i] \varphi_i \{..\} \rrbracket [s] \xrightarrow{\alpha} r$. Hence, knowing (30) and that $\llbracket \varphi' \rrbracket = \llbracket \bigwedge_{i \in I} [\alpha_i] \varphi_i \{..\} \rrbracket$, from (29) and (30) we conclude that $\llbracket \max X. \varphi' \rrbracket [s] \xrightarrow{\alpha} r$ as required, and so we are done. \square

Having concluded the proof of Theorem 2 and knowing Theorem 1, we can finally obtain our main result with respect to Definition 3.

Theorem 3 *Controlled system synthesis is the static counterpart of suppression enforcement in the context of safety properties.* \square

5 Distinguishing between Suppression Enforcement and CSS

Despite concluding that CSS is the static counterpart to suppression enforcement, there are still a number of subtle differences between these two techniques. For one, since suppression enforcement is a dynamic technique, the monitor and the system still remain two separate entities, and the instrumentation between them is merely a way for the monitor to interact with the SuS. In general, the monitor does not affect the execution of the SuS itself, but rather modifies its observable trace of actions, such as its inputs and outputs. By contrast, when a controlled system is synthesised, an existing system is paired up with a formula and statically reconstructed into a *new* (correct) system that is incapable of executing the erroneous behaviour.

By removing invalid transitions entirely, controlled system synthesis is more ideal to guarantee the property compliance of the *internal* (less observable) behaviour of a system. For example, this can be useful to ensure that the system does not use a shared resource before locking it. By contrast, the invalid actions are still executed by the system in suppression enforcement, but their effect is rendered invisible to any external observer. This makes suppression enforcement more suitable to ensure that the *external* (observable) behaviour of the system complies with a desired property. For instance, one can ensure that the system does not perform an output that is

innocuous to the system itself, but may be providing harmful information to the external environment.

Moreover, it turns out that although both techniques produce composite systems that are trace equivalent to each other, an external observer may still be able to tell them apart by merely observing them. One way of formally assessing this is to use observational equivalence (characterised by weak bisimilarity) as a yardstick, thus:

$$\forall \varphi \in \text{SHML}, s \in \text{SYS}, \exists m \in \text{TRN} \cdot m[s] \approx (s, \varphi). \quad (32)$$

We show by means of a counter example that (32) is in fact *false* and as a result prove Theorem 4.

Theorem 4 (Observational Difference) *There exist an SHML_{inv} formula φ , an LTS $\langle \text{SYS}, \text{ACT}, \rightarrow \rangle$ and a system state $s \in \text{SYS}$ such that for all monitors $m \in \text{TRN}$, $m[s] \not\approx (s, \varphi)$.* \square

Proof sketch. Recall the controlled LTS with initial state $(s_{\mathbf{b}}, \varphi_0)$ obtained in Example 4. To prove Theorem 4 we must show that for every action suppression monitor m (that can only apply suppression and identity transformations), one *cannot* find a weak bisimulation relation \mathcal{R} so that $(m[s_{\mathbf{b}}], (s_{\mathbf{b}}, \varphi_0)) \in \mathcal{R}$. An elegant way of showing this claim, is by playing the *weak bisimulation games* [3] starting from the pair $(m[s_{\mathbf{b}}], (s_{\mathbf{b}}, \varphi_0))$, for every possible m . The game is played between two players, namely, the attacker and the defender. The attacker wins the game by finding a sequence of moves from the monitored state $m[s_{\mathbf{b}}]$ (or the controlled state $(s_{\mathbf{b}}, \varphi_0)$), which the defender cannot counter, *i.e.*, the move sequence cannot be performed by the controlled state $(s_{\mathbf{b}}, \varphi_0)$ (*resp.* monitored state $m[s_{\mathbf{b}}]$). Note that the attacker is allowed to play a transition from either the current monitored state or the controlled state at each round of the game. A winning strategy for the attacker entails that the composite systems are *not* observationally equivalent.

We start playing the game from the initial pair $(m[s_{\mathbf{b}}], (s_{\mathbf{b}}, \varphi_0))$ for every monitor m . Pick any monitor that suppresses any action other than a second consecutive `ans`, such as $m_0 \stackrel{\text{def}}{=} \{\text{req}, \tau\}.m'_0$. In this case, it is easy to deduce that the defender always loses the game, that is, if the attacker attacks with $(s_{\mathbf{b}}, \varphi_0) \xrightarrow{\text{req}} (s_{\mathbf{b}}^1, \varphi_0)$ the defender is defenceless since $m_0[s_{\mathbf{b}}] \not\xrightarrow{\text{req}}$. This remains true regardless of the “depth” at which the suppression of the first `req` transition occurs.

On the one hand, using the same game characterisation, one can also deduce that by picking a monitor that *fails to suppress* the second consecutive `ans` action, such as $m_1 \stackrel{\text{def}}{=} \{\text{req}\}.\{\text{ans}\}.\{\text{ans}\}.m'_1$, also prevents the defender from winning. If the attacker plays with $m_1[s_{\mathbf{b}}] \xrightarrow{\text{req.ans.ans}} m'_1[s_{\mathbf{b}}]$, the defender loses since it can only counter the first two transitions, *i.e.*, $(s_{\mathbf{b}}, \varphi_0) \xrightarrow{\text{req.ans}}$

$\xrightarrow{\text{ans}}$. Again, this holds regardless of the “depth” of the first such failed suppression.

On the other hand, any monitor that actually suppresses the second consecutive `ans` action, such as $m_2 \stackrel{\text{def}}{=} \{\text{req}\}.\{\text{ans}\}.\{\text{ans}, \tau\}.m'_2$, still negates a win for the defender. In this case, the attacker can play $(s_{\mathbf{b}}, \varphi_0) \xrightarrow{\text{req.ans}} (s_{\mathbf{b}}^2, \varphi_0 \wedge [\text{ans}]ff)$ to which the defender must reply with $m_2[s_{\mathbf{b}}] \xrightarrow{\text{req.ans}} \{\text{ans}, \tau\}.m'_2[s_{\mathbf{b}}^2]$. The attacker can subsequently play $\{\text{ans}, \tau\}.m'_2[s_{\mathbf{b}}^2] \xrightarrow{\tau} m'_2[s_{\mathbf{b}}^1]$, which can only be countered by an inaction on behalf of the defender, *i.e.*, the controlled system remains in state $(s_{\mathbf{b}}^2, \varphi_0 \wedge [\text{ans}]ff)$.

Since we do not know the form of m'_2 , we consider the following two cases, namely, when m'_2 suppresses `cls`, and the case when it does not. In the first case, the attacker can attack with $m'_2[s_{\mathbf{b}}^1] \xrightarrow{\tau} m''_2[\text{nil}]$ (for some m''_2) where τ represents the suppression of the `cls` action. Once again the defender can only counter with an inaction and stay in state $(s_{\mathbf{b}}^2, \varphi_0 \wedge [\text{ans}]ff)$. At this point the attacker wins the play by attacking with $(s_{\mathbf{b}}^2, \varphi_0 \wedge [\text{ans}]ff) \xrightarrow{\text{log}} (s_{\mathbf{b}}^3, \varphi_0)$ since $m''_2[\text{nil}] \xrightarrow{\text{log}}$. In the second case, the attacker also wins by attacking with $m'_2[s_{\mathbf{b}}^1] \xrightarrow{\text{cls}} m'''_2[\text{nil}]$ (for some m'''_2) since $(s_{\mathbf{b}}^2, \varphi_0 \wedge [\text{ans}]ff) \xrightarrow{\text{cls}}$.

The same result can be obtained using monitor $m_{\mathbf{s}}$ from Example 2. In this case, the attacker can play $(s_{\mathbf{b}}, \varphi_0) \xrightarrow{\text{req.ans}} (s_{\mathbf{b}}^2, \varphi_0 \wedge [\text{ans}]ff)$ to which the defender can only reply with $m_{\mathbf{s}}[s_{\mathbf{b}}] \xrightarrow{\text{req.ans}} m'_s[s_{\mathbf{b}}^2]$. The attacker can subsequently play $m'_s[s_{\mathbf{b}}^2] \xrightarrow{\tau} \text{sup}[s_{\mathbf{b}}^1]$, which can only be countered by an inaction on behalf of the defender, *i.e.*, the controlled system remains in state $(s_{\mathbf{b}}^2, \varphi_0 \wedge [\text{ans}]ff)$. However, the attacker can subsequently play $(s_{\mathbf{b}}^2, \varphi_0 \wedge [\text{ans}]ff) \xrightarrow{\text{log}} (\text{nil}, \varphi_0)$ which is indefensible since $\text{sup}[s_{\mathbf{b}}] \xrightarrow{\text{log}}$.

These cases therefore suffice to deduce that for every possible monitor the attacker always manages to win the game, and hence we conclude that Theorem 4 holds as required. \square

This result is important since it proves that powerful external observers, such as the ones presented by Abramsky in [1], can still distinguish between the resulting monitored and controlled systems.

6 Related Work

Several works comparing formal verification techniques can be found in the literature. In [25] van Hulst *et al.* explore the relationship between their work on controlled system synthesis and the synthesis problem in Ramadge and Wonham’s Supervisory Control Theory (SCT) [32]. The aim in SCT is to generate a *supervisor controller*

from the SuS and its specification (*e.g.*, a formal property). If successfully generated, the synchronous product of the SuS and the controller is computed to obtain a supervised system. To enable the investigation, van Hulst *et al.* developed language-based notations akin to that used in [32], and proved that Ramadge and Wonham’s work can be expressed using their theory.

Ehlers *et al.* in [15] establish a connection between SCT and reactive synthesis – a formal method that attempts to automatically derive a valid reactive system from a given specification. To form this connection, the authors first equalise both fields by using a simplified version of the standard supervisory control problem and focus on a class of reactive synthesis problems that adhere to the requirements imposed by SCT. They then show that the supervisory control synthesis problem can be reduced to a reactive synthesis problem.

Basile *et al.* in [11] explore the gap between SCT and coordination of services, which describe how control and data exchanges are coordinated in distributed systems. This was achieved via a new notion of controllability that allows one to reduce the classical SCT synthesis algorithms to produce orchestrations and choreographies describing the coordination of services as contract automata.

Falcone *et al.* made a brief, comparison between runtime enforcement and SCT in [17] in the context of K-step opacity, but established no formal results that relate these two techniques.

7 Conclusion

We have presented a novel comparison between suppression enforcement and controlled system synthesis – two verification techniques that automate system correction for erroneous systems. We were able to conclude that controlled system synthesis is the static counterpart to suppression enforcement in the context of safety, as defined by Definition 3. This required developing a function that maps logic formulas to suppression monitors, Definition 5, and proving inductively that for every system and formula, one can obtain a monitored and a controlled system that execute the same set of traces at runtime, Theorem 2. As trace equivalent systems satisfy the same safety properties, Theorem 1, this result was enough to reach our conclusion, Theorem 3. Using a counter-example we however deduced that these two techniques are different modulo observational equivalence, Theorem 4. An Abramsky-type external observer [1] can therefore tell the difference between a monitored and controlled system resulting from the same formula and SuS. To our knowledge this is the first formal comparison to be made between these two techniques.

Future Work Having established a connection between suppression enforcement and control system synthesis with respect to safety properties, it is worth expanding this work at least along two directions and explore how:

- (i) runtime enforcement and controlled system synthesis are related with respect to properties other than those representing safety, and how
- (ii) suppression enforcement relates to other verification techniques such as supervisory control theory, reactive synthesis, *etc.*

Exploring (i) may entail looking into other work on enforcement and controlled system synthesis that explores a wider set of properties. It might be worth investigating how other enforcement transformations, such as action replacements and insertions, can be used to widen the set of enforceable properties, and how this relates to controlled system synthesis. The connection established by van Hulst *et al.* in [25] between control system synthesis and supervisory control, along with the other relationships reviewed in Section 6, may be a starting point for conducting our future investigations on (ii).

References

1. Abramsky S (1987) Observation equivalence as a testing equivalence. *Theoretical Computer Science* 53:225–241, URL [https://doi.org/10.1016/0304-3975\(87\)90065-X](https://doi.org/10.1016/0304-3975(87)90065-X)
2. Aceto L, Ingólfssdóttir A (1999) Testing hennessy-milner logic with recursion. In: Thomas W (ed) *Foundations of Software Science and Computation Structures*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 41–55
3. Aceto L, Ingólfssdóttir A, Larsen KG, Srba J (2007) *Reactive Systems: Modelling, Specification and Verification*. Cambridge University Press, New York, NY, USA
4. Aceto L, Achilleos A, Francalanza A, Ingólfssdóttir A, Kjartansson SÖ (2016) Determinizing monitors for HML with recursion. arXiv preprint
5. Aceto L, Achilleos A, Francalanza A, Ingólfssdóttir A (2018) A framework for parameterized monitorability. In: *Foundations of Software Science and Computation Structures*, Springer International Publishing, Cham, pp 203–220
6. Aceto L, Cassar I, Francalanza A, Ingólfssdóttir A (2018) On runtime enforcement via suppressions. In: *29th International Conference on Concurrency Theory, CONCUR 2018*, September 4–7, 2018, Beijing, China, pp 34:1–34:17, DOI 10.4230/LIPIcs.CONCUR.2018.34
7. Aceto L, Achilleos A, Francalanza A, Ingólfssdóttir A, Lehtinen K (2019) Adventures in monitorability: From branching to linear time and back again. *Proc ACM Program Lang* 3(POPL):52:1–52:29, DOI 10.1145/3290365, URL <http://doi.acm.org/10.1145/3290365>
8. Aceto L, Cassar I, Francalanza A, Ingólfssdóttir A (2019) Comparing controlled system synthesis and suppression enforcement. In: *Runtime Verification*, Springer International Publishing, Cham, pp 148–164
9. Alur R, Černý P (2011) Streaming transducers for algorithmic verification of single-pass list-processing programs. In: *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM, pp 599–610
10. Arnold A, Walukiewicz I (2008) Nondeterministic controllers of nondeterministic processes. In: *Logic and Automata*, Amsterdam University Press, *Texts in Logic and Games*, vol 2, pp 29–52
11. Basile D, ter Beek MH, Pugliese R (2019) Bridging the gap between supervisory control and coordination of services: Synthesis of orchestrations and choreographies. In: *COORDINATION 2019 - 21st International Conference on Coordination Models and Languages*, (To appear)
12. Cassar I, Francalanza A, Aceto L, Ingólfssdóttir A (2017) A survey of runtime monitoring instrumentation techniques. In: *PrePost2017*, pp 15–28
13. Clarke EM, Grumberg O, Peled D (1999) *Model Checking*. MIT press
14. Desai A, Dreossi T, Seshia SA (2017) Combining model checking and runtime verification for safe robotics. In: *Runtime Verification (RV)*, Springer International Publishing, Cham, LNCS, pp 172–189
15. Ehlers R, Lafortune S, Tripakis S, Vardi M (2014) Bridging the gap between supervisory control and reactive synthesis: Case of full observation and centralized control. *IFAC Proceedings Volumes* 47(2):222 – 227, 12th IFAC International Workshop on Discrete Event Systems (2014)
16. Erlingsson U, Schneider FB (1999) Sasi enforcement of security policies: A retrospective. In: *Proceedings of the 1999 Workshop on New Security Paradigms*, ACM, New York, NY, USA, NSPW '99, pp 87–95
17. Falcone Y, Marchand H (2013) Runtime enforcement of k-step opacity. In: *52nd IEEE Conference on Decision and Control*, pp 7271–7278, DOI 10.1109/CDC.2013.6761043
18. Falcone Y, Fernandez JC, Mounier L (2012) What can you verify and enforce at runtime? *International Journal on Software Tools for Technology Transfer* 14(3):349

19. Francalanza A (2016) A Theory of Monitors. In: International Conference on Foundations of Software Science and Computation Structures, Springer, pp 145–161
20. Francalanza A (2017) Consistently-Detecting Monitors. In: 28th International Conference on Concurrency Theory (CONCUR 2017), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, Leibniz International Proceedings in Informatics (LIPIcs), vol 85, pp 8:1–8:19
21. Francalanza A, Aceto L, Achilleos A, Attard DP, Cassar I, Della Monica D, Ingólfssdóttir A (2017) A foundation for runtime monitoring. In: Runtime Verification, Springer International Publishing, Cham, pp 8–29
22. Francalanza A, Aceto L, Ingólfssdóttir A (2017) Monitorability for the Hennessy-Milner logic with recursion. *Formal Methods in System Design* 51(1):87–116
23. Havelund K, Pressburger T (2000) Model checking java programs using java pathfinder. *International Journal on Software Tools for Technology Transfer* 2(4):366–381, DOI 10.1007/s100090050043, URL <https://doi.org/10.1007/s100090050043>
24. Havelund K, Roşu G (2004) An overview of the runtime verification tool java pathexplorer. *Formal methods in system design* 24(2):189–215
25. van Hulst AC, Reniers MA, Fokkink WJ (2017) Maximally permissive controlled system synthesis for non-determinism and modal logic. *Discrete Event Dynamic Systems* 27(1):109–142
26. Kejstová K, Ročkal P, Barnat J (2017) From Model Checking to Runtime Verification and Back. In: *RV*, Springer
27. Könighofer B, Alshiekh M, Bloem R, Humphrey L, Könighofer R, Topcu U, Wang C (2017) Shield synthesis. *Formal Methods in System Design* 51(2):332–361
28. Leucker M, Schallhart C (2009) A brief account of runtime verification. *The Journal of Logic and Algebraic Programming* 78(5):293–303
29. Ligatti J, Bauer L, Walker D (2005) Edit automata: enforcement mechanisms for run-time security policies. *International Journal of Information Security* 4(1):2–16
30. Milner R, Parrow J, Walker D (1992) A calculus of mobile processes, I. *Information and computation* 100(1):1–40
31. Pnueli A, Rosner R (1989) On the synthesis of a reactive module. In: *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM, New York, NY, USA, POPL '89, pp 179–190, DOI 10.1145/75277.75293, URL <http://doi.acm.org/10.1145/75277.75293>
32. Ramadge PJ, Wonham WM (1987) Supervisory control of a class of discrete event processes. *SIAM J Control Optim* 25(1):206–230
33. Sakarovitch J (2009) *Elements of Automata Theory*. Cambridge University Press, New York, NY, USA
34. Sangiorgi D (2011) *Introduction to Bisimulation and Coinduction*. Cambridge University Press, New York, NY, USA
35. Schneider FB (2000) Enforceable security policies. *ACM Transactions on Information and System Security (TISSEC)* 3(1):30–50