
Test selection for traces refinement

Ana Cavalcanti

*University of York, Department of Computer Science
York YO10 5GH, UK*

Marie-Claude Gaudel

*LRI, Université de Paris-Sud and CNRS
Orsay 91405, France*

Abstract

Theories for model-based testing identify exhaustive test sets: typically infinite sets of tests whose execution establishes the conformance relation of interest. Practical techniques rely on selection strategies to identify finite subsets of these tests, and popular approaches are based on requirements to cover the model. In previous work, we have defined testing theories for refinement-based process algebra, namely, CSP and *Circus*, a state-rich process algebra. In this paper, we consider the selection of tests designed to establish traces refinement. In this case, conformance does not require that all traces of the model are available in the system under test, and this can raise challenges regarding coverage criteria for selection. To address these difficulties, we present a framework for formalising a variety of selection strategies. We exemplify its use in the formalisation of a selection criterion based on coverage of process communications for integration testing. We consider models written in *Circus*, whose symbolic testing theory facilitates the definition of uniformity and regularity hypotheses based on data operations, but also imposes extra challenges for selection of concrete tests. Our results, however, are relevant for any formalism where the conformance relation does not require all the traces of the specification to be executable by the system under test.

Keywords: process algebra, CSP, *Circus*, synchronisation coverage, integration testing

1. Introduction

We address in this paper the issue of selecting tests when using an abstract specification for testing that a system behaves like one of its refinements. In particular, we are interested in models written using process algebra like CSP [51] or *Circus* [17], and in traces refinement in that context.

CSP is a well established notation that has been in use for more than twenty years. The availability of a powerful model checker has led to acceptance both in academia and industry. In the public domain, we have reports on applications in hardware and e-commerce [2, 26], for example. CSP has been combined with a number of data-modelling languages to define notations that cope with state-rich reactive systems [21, 36, 52, 49]; *Circus* is one of these combinations with Z [62]. *Circus* has been used for modelling and verification of control systems specified in Simulink [9, 38]. It is currently being used to verify aerospace applications, including virtualisation software by the US Naval Research Laboratory [23].

The traditional area of application for CSP and *Circus* is refinement-based development and verification, not testing, but whenever a CSP or a *Circus* model is available, the possibility of using it for testing is attractive, especially in industry. Moreover, when the correctness criterion for development is refinement, it is natural to adopt that same criterion for testing. We have previously developed testing theories for CSP [10] and *Circus* [13]. We base our discussions here on the theory for *Circus*, which extends and generalises that for CSP. In addition, it enables richer selection strategies based on data types (as well as traces).

The various notions of refinement adopted in the context of CSP and *Circus* require that the set of traces of an implementation is included in that of the specification; in particular, this is exactly the definition of traces refinement. Selection of tests for traces refinement is challenging because it is acceptable that some traces of the specification are not implemented by the system under test (SUT). It may, therefore, be the case that some selected tests are not executable. As a consequence, given a selection criterion, the selection cannot be performed statically only. At runtime, the test driver must ensure that alternative tests are submitted until the criterion is fulfilled, or until it is certain that it cannot, due to lack of traces necessary to its fulfillment. In other words, any testing strategy requires a combination of some selection criterion with a driver designed to ensure that the selection criterion will be reached at run time, if it is possible.

Our main contributions are as follows. First, we propose a formal framework for selection based on the definition of an exhaustive set of tests that can contribute to the fulfillment of the selection criterion and of a test driver. Additionally, we consider symbolic tests, which pose a challenge of the same nature for obtaining concrete tests that can be used by the driver, since it is acceptable that some or all instantiations are not implemented. Symbolic tests lead to a two-level selection method, where first some symbolic tests are selected, and then we select some of their instantiations, which are submitted until the criterion is fulfilled. Finally, as an application of our new framework, we consider coverage of communications between parallel processes of a *Circus* model; we present a complete formalisation of this selection criterion.

Our work bears some similarity with the dynamic test generation techniques [6], where symbolic execution is used at runtime to ensure coverage of all feasible paths of a program text. Here, however, the issue is not infeasibility, but arbitrary partial implementation of traces. Moreover, we are concerned with selection criteria that are based on requirements to exercise elements of the model, not the SUT. Our framework also has some similarities with adaptive testing [33], which is aimed at tackling nondeterminism in implementations rather than partial implementation of traces. In contrast, for coverage of behaviours of a nondeterministic SUT, we rely on the complete testing assumption [24], which may, in special cases, be enforced by techniques similar to those used in reachability testing [34, 39].

In the next section, we give an overview of *Circus* and its testing theory. Section 3 presents our novel framework for selection. In Section 4, we formalise a synchronisation coverage criterion for use with our framework, and we define finite test sets and suitable drivers. We consider related works in Section 5, and conclude in Section 6 with a discussion of our results and future work.

2. *Circus* and its testing theory

In this section, we cover the background material for our work: we describe *Circus* using an example of a simple protocol, and introduce the main concepts of its testing theory.

2.1. An overview of *Circus*

In Figure 1, we present a *Circus* model for a simple protocol for communication between a *Sender* and a *Receiver* via a *Medium* that can corrupt at most one out of three messages. Systems and their components are modelled in *Circus* using processes. These encapsulate some state and interact with each other and their environment via channel communications, which are synchronous, atomic, and instantaneous events. In Figure 1, we have a model that uses four channels, *in*, *out*, *left*, and *right*, and introduces five processes: *Medium*, *Sender*, *Receiver*, *Components*, and *Protocol*.

We observe that, although *Circus* events (channel communications) are synchronous, this example shows that its modelling paradigm is rich enough to cater for asynchronous systems. Our example models an asynchronous communication between the *Sender* and the *Receiver* by identifying the relevant interaction points of the asynchronous transaction (using the channels *left* and *right*).

A *Circus* model is formed by a sequence of paragraphs. In Figure 1, the first paragraph defines a set *Bit* including 0 and 1. The messages exchanged using the specified protocol are bits. In the second paragraph, we declare the channels, whose type *Bit* defines that they are used to communicate bits.

The next three paragraphs in Figure 1 introduce (basic) processes by (explicitly) defining their states and behaviour. Process states are defined using Z schemas. The state of the process *Medium* has a single

```

Bit == {0, 1}

channel in, out, left, right : Bit

process Medium ≐ begin
  state MS == [ noE : 0 .. 2 ]
  • noE := 0 ; ( μ X • ( left?x → ( if noE = 0 →
    ( right!x → Skip □ right!(1 - x) → noE := 2 )
    □ noE > 0 → right!x → noE := noE - 1
    fi
  ) ) ; X )
end

process Sender ≐ begin
  • ( μ X • in?m → left!m → left!m → left!m → X )
end

process Receiver ≐ begin
  state MS == [ ms : bag Bit ]
  • ( μ X • ( ms := [ ] ; right?m1 → right?m2 → right?m3 → )
    ms := [ m1, m2, m3 ] ; out!(maj ms) → X )
end

process Components ≐ Sender ||| Receiver

process Protocol ≐ ( Components [ { left, right } ] Medium ) \ { left, right }

```

Figure 1: *Circus* example: a simple protocol

component *noE* recording the number of messages that must be sent correctly before another error is possible. We use this information to model the assumption that at most one out of every three messages is corrupted. The declaration of *noE* indicates that it can only take the values 0, 1, and 2.

After the •, we have a *Circus* action that specifies the behaviour of *Medium*. It is written using a combination of CSP operators and guarded commands. First of all, an assignment initialises *noE* to 0, so that a corrupt message can be sent immediately. Afterwards, a recursive action iteratively accepts messages through the channel *left* and sends them through *right*. Recursion is defined using the operator $\mu X \bullet A(X)$, which introduces the local name *X* for the action *A*, in which references to *X* are recursive calls.

In the body of the recursion in *Medium*, we have an input communication *left?**x* that accepts an input through the channel *left* and records it in the local variable *x*. The prefixing operator $c \rightarrow A$ describes a behaviour in which, after the communication *c*, we have the execution of the action *A*. In *Medium*, after the communication *left?**x*, we have a conditional that checks the value of *noE*. If it is 0, then it is possible that a corrupt message is sent. In this case, we have a nondeterministic choice between outputting through the channel *right* the input *x* (communication *right!**x*) or the corrupt value $1 - x$, in which case the value 2 is assigned to *noE*. The nondeterministic choice operator $A_1 \sqcap A_2$ chooses arbitrarily between the execution of *A*₁ or *A*₂. In our example, the nondeterminism reflects the fact that we are abstracting from the possible causes of a medium error, so that the occurrence of an error becomes completely arbitrary.

If the action chosen is *right!**x* → **Skip**, then after the output over *right*, it behaves like **Skip**, which terminates immediately without changing the state. If the other action is chosen, after the output, the value of *noE* is updated to indicate that the next two communications must not be corrupt. In the conditional, if the value of *noE* is not 0, then *right!**x* takes place, and the value of *noE* is decremented.

The process *Sender* has no state. Its action defines that it accepts an input *m* through the channel *in* and sends it three times to the *Medium* through the channel *left*, before recursing and starting again.

The *Receiver* process, on the other hand, keeps a bag of bits *ms* in the state. At each iteration of the

recursion in its action, the bag is emptied ($ms := \llbracket \ \rrbracket$), three messages m_1 , m_2 , and m_3 are accepted (from the *Medium*) through *right*. These messages are stored in ms , and then the bit ($maj\ ms$) with the majority of occurrences in ms is output through *out*. (The definition of $maj\ ms$ is omitted in Figure 1.)

Circus processes can also be combined using CSP operators. In our example, the components connected by the protocol are modelled by the process *Components*, defined by the interleaving of the processes *Sender* and *Receiver*. This reflects the fact that these processes are executed in parallel and independently. The *Protocol* process itself is defined in terms of the parallel composition of *Components* and *Medium*. In this case, we have a parallelism, rather than an interleaving, because these processes need to synchronise on communications over the channels *left* and *right*. In a parallelism $P_1 \llbracket cs \rrbracket P_2$, the processes P_1 and P_2 proceed independently, except only for communications on channels in the given set cs , on which they are required to synchronise. In particular, the state of P_1 is not visible to P_2 , and vice-versa.

In the definition of *Protocol*, we also use hiding. In a process $P \setminus cs$ the communications over the channels in the set cs are hidden: they are internal and carried out as soon as possible. In our example, *left* and *right* are channels used by the *Medium* to connect the *Components*, and so they are internal. Communication with the external environment is realised via the channels *in* (used in *Sender*) and *out* (used in *Receiver*).

We use this and other examples to illustrate our results. Any extra use of *Circus* notation is explained.

The semantics of *Circus* is defined in a relational style using Hoare and He's Unifying Theories of Programming (UTP) [43]; operational semantics for *Circus* are defined in [13, 12]. In [12] we define traces refinement using the UTP model. As expected, the definition compares processes P_1 and P_2 with the same alphabet of events, and requires that P_1 is refined by P_2 if, and only if, every trace of P_2 is also a trace of P_1 . In the next section, we give an overview of the *Circus* theory for testing for traces refinement.

2.2. Testing against traces refinement using *Circus* (and CSP)

In a theory for model-based testing for a particular modelling notation and conformance relation, we identify testability hypotheses, notions of test, test execution, and verdict, and an exhaustive test set. This is a set of tests whose execution and associated verdicts can determine unequivocally whether the SUT behaves in conformance to a given model, as long as the testing hypotheses hold. Typically, this is an infinite set of tests. Here, we consider selection as the definition of a subset of the exhaustive test set. Thus, we briefly recall the definition of the exhaustive test set for *Circus* and traces refinement [13].

The two testability hypotheses of the *Circus* theory are standard. First, we assume that the SUT can be described by an (unknown) *Circus* model. This means, for example, that events are atomic (indivisible) and can be considered as instantaneous: their duration can be neglected when observed at the level of the system. Secondly, we rely on the complete test assumption [24]. It is concerned with possible nondeterminism in the SUT (as opposed to its model); it requires the existence of a number n such that, if a test is executed n times, then all possible (nondeterministic) behaviours of the SUT are observed. For a deterministic SUT, n is 1. When the only source of nondeterminism is due to interleaving of concurrent sequential executions, the complete test assumption can be enforced by the use of reachability-testing techniques to monitor the executions of the SUT. Some of these techniques are briefly presented in Section 5.

The tests of the *Circus* theory embed the verdict. We use special verdict events *inc*, *pass*, and *fail* used to indicate that a test is *inconclusive*, because the SUT did not execute the particular trace of events attempted by the test, or that the SUT *passed* or *failed* the test. A test is constructed based on a particular trace t of the model, and a forbidden continuation e of t . A forbidden continuation is an event that is not allowed after the trace t ; in other words, the trace $t \hat{\ } \langle e \rangle$ is not a trace of the model. (We use $\langle e \rangle$ to denote the singleton sequence containing just the event e , and $\hat{\ }$ is the concatenation operator.)

Example 1. *In our example, $\langle in.0, out.0, in.1 \rangle$ is a trace of *Protocol*. The event $in.0$, for instance, corresponds to a value 0 input through the channel *in*, and, similarly, $out.0$, the value 0 output through the channel *out*. After the trace $\langle in.0, out.0, in.1 \rangle$, all events $out.x$, where $x \neq 1$ are forbidden continuations of *Protocol*, because after an input of a value 1, just 1 itself can be output. \square*

A test for a trace t and a forbidden continuation e attempts to drive the SUT to execute t . If it does not succeed, the test is inconclusive. (As usual, we give an inconclusive verdict when the intended test is

not completed.) If the test does succeed in driving the SUT to execute t , then the forbidden event e is attempted, and if it is accepted, the SUT fails the test. Otherwise, it succeeds.

Example 2. For the trace in Example 1, and forbidden continuation *out.2*, we have the test below.

$$inc \longrightarrow in.0 \longrightarrow inc \longrightarrow out.0 \longrightarrow inc \longrightarrow in.1 \longrightarrow pass \longrightarrow out.2 \longrightarrow fail \longrightarrow \mathbf{Stop}$$

The action **Stop** deadlocks immediately. □

The verdict events are interspersed with the events of the trace and the forbidden continuation. The last verdict event before a deadlock determines the verdict of the test. Formally, tests are defined as follows.

Definition 1.

$$T_T(\langle \rangle, e) = pass \rightarrow e \rightarrow fail \rightarrow \mathbf{Stop} \qquad T_T(\langle e_1 \rangle \frown t, e_2) = inc \rightarrow e_1 \rightarrow T_T(t, e_2)$$

In $T_T(t, e)$, extending the trace t with e is supposed to lead to an invalid trace, since e is supposed to be a forbidden continuation: not an initial after t ; as already explained, the test aims at ruling it out.

Execution $Execution_{SUT}^{SP}(T)$ of one of these tests T is carried out by executing T and the SUT in parallel, synchronising on all model events (that is, the non-verdict events used in the model and in the SUT). Relying on our first testability hypothesis, we use SUT to denote the unknown model of the SUT. We also use Σ to denote the set of model events as defined in the specification SP .

Definition 2.

$$Execution_{SUT}^{SP}(T) = (SUT \llbracket \Sigma \rrbracket T) \setminus \Sigma \text{ where } \Sigma = \alpha SP$$

The model events are hidden, that is, internal to the experimentation process. This ensures that they are not affected by the environment of the experiment, since they are not visible in the environment. Hiding also ensures that these model events occur as soon as they become available (in the SUT and in the test).

For a given model SP , the exhaustive test set $Exhaust_T(SP)$ for traces refinement is built by considering all traces of SP and all their forbidden continuations. We consider the set $traces(SP)$ [11], which characterises a trace semantics for SP . We also consider the set $initials(SP, t)$ of continuations of the trace t of SP .

Definition 3 (Exhaustive test set of SP for traces refinement). Given a specification SP , we define $Exhaust_T(SP) = \{ T_T(t, e) \mid t \in traces(SP) \wedge e \notin initials(SP, t) \}$.

Exhaustiveness of $Exhaust_T(SP)$ is formally established by the following theorem, proved in [10].

Theorem 1 (Exhaustivity of $Exhaust_T$). Given two processes, SUT and SP , $SP \sqsubseteq_T SUT$ if, and only if $\forall T_T(s, a) : Exhaust_T(SP); t : traces(Execution_{SUT}^{SP}(T_T(s, a))) \bullet last(t) \neq fail$

This extreme test strategy requires the execution of all the traces of SP with all their forbidden continuations. It yields an inconclusive verdict when a tested trace is not implemented. If there are only *pass* and *inc* verdicts, it guarantees that the SUT refines SP providing our testability hypotheses hold.

As said above, when selecting a strict subset of $Exhaust_T$, and running the tests T of this set as described by $Execution_{SUT}^{SP}(T)$ an element to be covered can be statically covered in the subset, but not covered at runtime. In Section 3, we introduce the notions of exhaustive coverage of an element of the specification, and runtime coverage of this element by the implementation. Before discussing these novel concepts, however, we discuss briefly a symbolic account of the *Circus* tests.

2.3. Symbolic tests in Circus

The testing theory described above was first cast in the context of CSP [10], rather than *Circus*. The main difference between *Circus* and CSP, however, is the added possibility of specifying complex data types using the predicative style of pre and postconditions. Since *Circus* processes encapsulate the data, it is not surprising that tests can be constructed from a *Circus* process in the same way as from a CSP process. With similar notions of traces and trace refinement, all results of the CSP theory are maintained.

On the other hand, it is possible and, of interest, to take into account the data operations of a model when specifying a testing strategy. To this end, we have developed a symbolic characterisation of tests and exhaustive test sets for *Circus* [13]. It is based on the *Circus* operational semantics, which maintains a symbolic account of the state of a process. We give a brief description of the main definitions here.

The *Circus* operational semantics defines a transition relation $p_1 \xrightarrow{1} p_2$ between texts p_1 and p_2 of processes. When $p_1 \xrightarrow{1} p_2$ holds, we say that there is a transition from p_1 to p_2 with label 1. We use the typewriter font to distinguish texts (of processes, actions, predicates, and so on), from the entities that they denote. If the label is ϵ , the execution of the process p_1 can evolve silently (without any interaction with the environment), so that we then have an execution of p_2 . On the other hand, there may be a label $d.w$, representing a communication over the channel d of a value represented by the symbolic variable w . In this case, the execution of p_1 evolves to that of p_2 , after engaging in the communication $d.w$.

The definition of this transition relation for processes is based on a transition relation for configurations involving actions. They are triples $(c \mid s \models A)$ where c is a constraint over the symbolic variables in use, s is an assignment to the *Circus* variables in scope of values represented by symbolic variables, and A is the text of a *Circus* action. The symbolic variables constrained in c are those used in labels $d.w$ to represent communications and those used in the definition of the state s to represent values of *Circus* variables.

To construct symbolic tests, we use the notion of constrained symbolic traces of a process. A constrained symbolic trace is a pair formed by a symbolic trace st and a constraint c . A symbolic trace is a finite sequence of symbolic events $d.\alpha_0$, where d is a channel, and α_0 is a symbolic variable that represents the value communicated. The constraint c is the text of a predicate over the symbolic variables used in st .

Example 3. For the *Protocol* process, we have the constrained symbolic trace $(\langle in.\alpha, out.\beta, in.\gamma \rangle, \alpha = \beta)$, for example. It represents all traces where the first input and output values α and β are the same, and then a new, unconstrained value γ is input. The trace in Example 1 is one of its instances. \square

We define below the function $cstraces$ that characterises the set of constrained symbolic traces of a process. It is parametrised by an alphabet a , which fixes the symbolic variables that can be used in the symbolic trace, and the order in which they can be used. For a process `begin state $[x : T] \bullet A$ end`, the constrained symbolic traces are those of its main action A , starting from a state in which the state x takes the value w_0 constrained by $w_0 \in T$. This is the set $cstraces^a(w_0 \in T, x := w_0, A)$ defined using the operational semantics.

Definition 4.

$$cstraces^a(\text{begin state}[x : T] \bullet A \text{ end}) = cstraces^a(w_0 \in T, x := w_0, A)$$

$$cstraces^a(c_1, s_1, A_1) = \left\{ \begin{array}{l} st, c_2, s_2, A_2 \mid \alpha st \leq a \wedge (c_1 \mid s_1 \models A_1) \xrightarrow{st} (c_2 \mid s_2 \models A_2) \\ \bullet \mathcal{R}(st, c_2) \end{array} \right\}$$

In defining $cstraces^a(c_1, s_1, A_1)$, we consider the configurations $(c_2 \mid s_2 \models A_2)$ that can be reached by evaluation of the operational semantics from the configuration $(c_1 \mid s_1 \models A_1)$ identified by the parameters. They are characterised by the transition relation $(c_1 \mid s_1 \models A_1) \xrightarrow{st} (c_2 \mid s_2 \models A_2)$, which is annotated with a trace st . This transition relation is defined in terms of the operational semantics in the usual way.

The symbolic trace st records the events that arise from the evaluation of $(c_1 \mid s_1 \models A_1)$ in reaching $(c_2 \mid s_2 \models A_2)$ as illustrated in Example 3. With $\alpha st \leq a$, we require that the symbolic variables used to define st , like α, β and γ in Example 3, are those of the alphabet a , in the order in which they appear in a . We use αst to refer to the sequence of variables used in st , and $\alpha st \leq a$ requires that it is a prefix of a .

The constraint c_2 associated with \mathbf{st} , however, includes constraints on symbolic variables that are not used in \mathbf{st} , but are used to represent internal values arising from the evaluation of $(c_1 \mid \mathbf{s}_1 \models \mathbf{A}_1)$. These symbolic variables should not be mentioned in the constrained symbolic trace. Above, we use the function \mathcal{R} defined below to remove them. When applied to a trace $(\mathbf{st}, \mathbf{c})$, it quantifies away every variable in the set $\alpha\mathbf{c}$ of free variables of \mathbf{c} that is not in the set $\alpha\mathbf{st}$ of symbolic variables used in \mathbf{st} .

Definition 5. For every $(\mathbf{st}, \mathbf{c})$, we define $\mathcal{R}(\mathbf{st}, \mathbf{c})$ as follows.

$$\mathcal{R}(\mathbf{st}, \mathbf{c}) = (\mathbf{st}, \exists(\alpha\mathbf{c} \setminus \alpha\mathbf{st}) \bullet \mathbf{c})$$

The operational semantics guarantees that all configurations have satisfiable constraints. Therefore, every

constrained symbolic trace is feasible, in the sense that it has at least one valid instance, and so it is guaranteed that the constraint c_2 to which \mathcal{R} is applied in Definition 4 is satisfiable.

As said above, to construct tests, we also need forbidden continuations. In the symbolic approach, we use symbolic events: pairs formed by a symbolic event and a constraint. A symbolic event is itself formed by a channel name and a symbolic variable. We define the set $\overline{csinitials}^{\mathbf{a}}(\mathbf{P}, (\mathbf{st}, c))$ of constrained symbolic events over the alphabet \mathbf{a} that represent the events that are not initials of \mathbf{P} for any of the instances of its symbolic trace (\mathbf{st}, c) . A definition of this set is in [13], and its use for test generation is detailed in [20].

Example 4. The set $\overline{csinitials}^{\mathbf{a}}(\text{Protocol}, (\langle \text{in}.\alpha, \text{out}.\beta, \text{in}.\gamma \rangle, \alpha = \beta))$ of constrained symbolic events for Protocol after the trace in Example 3, where \mathbf{a} is the alphabet $\langle \alpha, \beta, \gamma, \delta, \dots \rangle$, contains the symbolic event $(\text{out}.\delta, \alpha = \beta \wedge \gamma \neq \delta)$. It determines that the value δ output after an input $\text{in}.\gamma$ is different from γ . We observe that the constraint is over the variable used in the symbolic event, as well as the variables used in the symbolic trace that lead to it. In this example, the constraint refers to α , β , and γ , as well as δ . \square

From the symbolic traces and events, we can construct symbolic tests in the expected way. For a constrained symbolic trace (\mathbf{st}, c_1) and a symbolic event $(d.\beta_0, c_2)$ that represents a forbidden continuation, the corresponding symbolic test is given by $\text{ST}_{\top}^{\alpha}((\mathbf{st}, c_1), (d.\beta_0, c_2))$ defined below.

Definition 6.

$$\begin{aligned} \text{ST}_{\top}^{\alpha}(\langle \rangle, c_1), (d.\beta_0, c_2) &= \text{pass} \longrightarrow d?\beta_0 : c_2 \longrightarrow \text{fail} \longrightarrow \text{Stop} \\ \text{ST}_{\top}^{\alpha}(\langle d_{\text{st}}.\alpha_0 \rangle \wedge \mathbf{st}, c_1), (d.\beta_0, c_2) &= \text{inc} \longrightarrow d_{\text{st}}?\alpha_0 : (\exists \overline{\alpha}, \overline{\alpha_0} \bullet c_1) \longrightarrow \text{ST}_{\top}^{(\alpha, \alpha_0)}((\mathbf{st}, c_1), (d.\beta_0, c_2)) \end{aligned}$$

For a constrained symbolic trace $(\langle \rangle, c_1)$ whose symbolic trace is empty, we only add to the test the verdict event *pass*. After that, corresponding to a symbolic forbidden continuation $(d.\beta_0, c_2)$, we have an input communication $d?\beta_0 : c_2$, which accepts an input β_0 through the channel d , with the constraint c_2 restricting the accepted values of the variable β_0 . Similarly, for the prefixings corresponding to each of the symbolic events $d_{\text{st}}.\alpha_0$ of the symbolic trace \mathbf{st} , we use the constraint c_1 to extract the right constraint over the input α_0 . The α parameter of ST_{\top} records the symbolic variables already used, and the constraint on α_0 is obtained by quantifying all variables not in α and different from α_0 , namely, those in the list $\overline{\alpha}, \overline{\alpha_0}$.

The instantiation of a symbolic test is defined below as a function that yields a set of concrete tests.

Definition 7.

$$\begin{aligned} \text{instTest}(\text{Stop}) &= \{ \text{Stop} \} \\ \text{instTest}(d \longrightarrow \mathbf{A}) &= \{ TA : \text{instTest}(\mathbf{A}) \bullet d \longrightarrow TA \} \\ \text{instTest}(d?\alpha_0 : c \longrightarrow \mathbf{A}) &= \{ v_0, TA \mid c[v_0/\alpha_0] \wedge TA \in \text{instTest}(\mathbf{A}[v_0/\alpha_0]) \bullet d.v_0 \longrightarrow TA \} \end{aligned}$$

The instances of a symbolic test are built in the following way. The action *Stop* is not really symbolic; its only instance is itself. For a prefixing $d \longrightarrow \mathbf{A}$, whose communication is a simple synchronisation like *inc*, *pass*, or *fail*, for example, the instances are prefixings formed out of d itself and the instances TA of the prefixed action \mathbf{A} . Finally, if we have a prefixing $d?\alpha_0 : c \longrightarrow \mathbf{A}$, then the instances consider all possible ways of choosing for communication a value v_0 that satisfies the constraint c , and the corresponding instances of the action $\mathbf{A}[v_0/\alpha_0]$, where the choice of v_0 for α_0 is recorded.

Example 5. For the constrained symbolic trace and forbidden continuation in Example 4, we get the following symbolic test.

$$\begin{aligned} \text{inc} \longrightarrow \text{in}.\alpha : \text{true} \longrightarrow \text{inc} \longrightarrow \text{out}.\beta : \alpha = \beta \longrightarrow \text{inc} \longrightarrow \text{in}.\gamma : \alpha = \beta \longrightarrow \\ \text{pass} \longrightarrow \text{out}.\delta : \alpha = \beta \wedge \gamma \neq \delta \longrightarrow \text{fail} \longrightarrow \text{Stop} \end{aligned}$$

The test in Example 2 is an instance of this test. \square

A slightly different definition of instantiation that avoids the enumeration of all the instances of the final forbidden symbolic events is possible: given a symbolic test constructed from the constrained symbolic trace (st, c_1) followed by the forbidden constrained symbolic event $(d.\beta_0, c_2)$, the variables in (st, c_1) are instantiated with values that satisfy the constraint c_1 , and these instantiations are propagated to c_2 as described above, but the final constraint is not instantiated. In this case, instances of symbolic tests are *Circus* processes where the single constrained symbolic event in the symbolic test is used in a final synchronisation with a predicate. Such an instance of the symbolic test of Example 5 is as follows.

$$\text{inc} \longrightarrow \text{in}.0 \longrightarrow \text{inc} \longrightarrow \text{out}.0 \longrightarrow \text{inc} \longrightarrow \text{in}.1 \longrightarrow \text{pass} \longrightarrow \text{out}.\delta : 1 \neq \delta \longrightarrow \text{fail} \longrightarrow \text{Stop}$$

A symbolic exhaustive test set is just a set of symbolic tests. Selection strategies based on symbolic tests are the subject of Section 3, selection of instances is addressed in Section 3.4, and further selection focused on coverage of synchronisations in *Circus* model is studied in Section 4.5.

3. Selection and traces refinement

The criteria for selecting a subset of the generally infinite set of possible tests are the core of any testing strategy. This can be achieved via coverage criteria, test purposes, uniformity or regularity hypotheses, or more generally some property of tests [28]. In the testing theory for CSP and *Circus*, tests are built from traces of the specification and a forbidden continuation; the exhaustive test set covers all these traces and events. So, selection must ultimately define a subset of the set of traces and their forbidden continuations.

As already mentioned, when the considered conformance relation is traces refinement, it is acceptable that a trace of the specification is not implemented in the SUT. As seen in Section 2.2, in this case execution of the corresponding test yields an *inconclusive* verdict. Thus it may be the case that a selected test is not executable, and if there is another test that may be able to fulfill the selection criterion, it must be executed, and we should proceed in this way, until either a satisfactory test is executed or it is certain that none exists. This has an impact on the way in which test selection for traces refinement can be defined and applied.

There are two main interrelated tasks involved in the definition of a selection strategy for traces refinement: a (formal) definition of (a) the (possibly still infinite) subset of all those tests that can contribute to the fulfillment of the selection criterion and (b) a test driver. We explain in the next section why defining an exhaustive test set with respect to the selection criteria is useful when considering traces refinement.

Whether concrete or symbolic tests are considered is irrelevant for definition of the exhaustive test set. We just observe that the notion of constrained symbolic traces, which is central in the *Circus* testing theory, is well suited for the definition of uniformity subdomains in the selection of both traces and initials: the constraints of the symbolic traces and events can be reused. On the other hand, we need to construct a driver of concrete tests. When selection is based on symbolic tests, we have, in addition, to address the fact that it may be the case that not all instantiations of a selected symbolic test are executable by the SUT.

Next, we illustrate the problem of selection for traces refinement. In Section 3.2, we consider several approaches to the definition of an exhaustive test set and of a test driver with respect to various selection criteria. In Section 3.3, we address factorisation of tests as for adaptive testing [33], and briefly discuss online and offline test generation. Section 3.4 discusses symbolic tests. Finally, Section 3.5 considers the issues related to the finiteness of the exhaustive test sets with respect to the given criteria.

3.1. The problem

To explain the problem associated with test selection for traces refinement in detail, we consider the case where the selection is based on a coverage criterion, that is, a set of elements of the specification are required to be exercised by at least one test experiment. The elements may be, for example, synchronisation events, pairs of related events, or more sophisticated requirements on execution traces.

If the test set is selected in a static way, it may turn out that the test selected to ensure coverage of a given element reaches an *inconclusive* verdict before executing it because the trace is not implemented. It may be the case, however, that other traces that provide coverage are implemented and that covering it is achievable in the SUT. Thus it is necessary to select and design tests that try at runtime other ways of covering this element, in order to ensure that if it is executable by the SUT, then it is exercised.

Example 6. For simplicity, we consider channels a, b, c, d and e that are used for synchronisation, but do not communicate any value. In this case, the constrained symbolic traces degenerate to standard traces, because there is no need to introduce symbolic variables to represent communicated values, and the only valid constraint is just `true`. In addition, we consider a stateless *Circus* process P defined below.

$$P \hat{=} \mathbf{begin} \bullet a \rightarrow e \rightarrow c \rightarrow \mathbf{Skip} \sqcap b \rightarrow e \rightarrow d \rightarrow \mathbf{Skip} \mathbf{end}$$

In what follows, whenever we present a stateless process like this, for brevity, we omit the `begin` and `end` keywords and identify the process with its action after the ‘ \bullet ’.

Some of the processes that trace-refine P are P itself and processes whose main actions are either of the actions in the internal choice. A test set that covers the event e , for instance, is the singleton including the process whose action is `inc` \rightarrow `a` \rightarrow `inc` \rightarrow `e` \rightarrow `pass` \rightarrow `d` \rightarrow `fail` \rightarrow `Stop`. To be sure, however, that if a trace covering e exists in the SUT, then e is covered by some test experiment, we need to identify the set of all tests that provide the required coverage and use a new way of driving these tests. Alternatively, a new form of test may be used, which adapts at runtime to the SUT.

The first solution is to design a driver based on a set of all tests covering e that will trigger their executions until e is exercised. The second solution is to have tree-shaped tests such as:

$$\mathbf{inc} \rightarrow \left(\begin{array}{l} \mathbf{a} \rightarrow \mathbf{inc} \rightarrow \mathbf{e} \rightarrow \mathbf{pass} \rightarrow \mathbf{d} \rightarrow \mathbf{fail} \rightarrow \mathbf{Stop} \\ \square \\ \mathbf{b} \rightarrow \mathbf{inc} \rightarrow \mathbf{e} \rightarrow \mathbf{pass} \rightarrow \mathbf{c} \rightarrow \mathbf{fail} \rightarrow \mathbf{Stop} \end{array} \right)$$

This test offers to the SUT the choice between two traces of the specification that pass through e . With this test, if the SUT is deterministic, if there exists in the SUT a specified trace exercising e , it is executed. When the SUT is not deterministic, by applying the test a sufficient number of times, under the complete test assumption, we know that the trace is executed. If there is no such trace, the test is inconclusive. \square

In the next section we explore both solutions indicated in the above example.

3.2. A solution: monitoring runtime coverage

As said above, when combining test selection and traces refinement, use of a selection criterion requires both the selection of traces and forbidden continuations and the design of a driver that ensures that the criterion is fulfilled at runtime, if possible. We proceed by selecting all the tests that satisfy the criterion, and define the corresponding driver that tries all these tests if necessary until a test is successfully completed.

Example 7. For illustration, we consider as selection criterion coverage of a particular event e . In this case, given a specification SP , the exhaustive set of tests that satisfy the selection criterion is as follows.

$$Exhaust_T(SP)|_e = \{ T_T(t, a) \mid t \in traces(SP) \wedge t|_e \neq \langle \rangle \wedge a \notin initials(SP, t) \}$$

where $t|_e$ eliminates from t all events different from e . This subset of $Exhaust_T(SP)$ covers all the traces where e occurs at least once. For Example 6, the tests in this set are as follows.

```

a  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  pass  $\rightarrow$  a  $\rightarrow$  fail  $\rightarrow$  Stop
a  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  pass  $\rightarrow$  b  $\rightarrow$  fail  $\rightarrow$  Stop
a  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  pass  $\rightarrow$  d  $\rightarrow$  fail  $\rightarrow$  Stop
a  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  inc  $\rightarrow$  c  $\rightarrow$  pass  $\rightarrow$  a  $\rightarrow$  fail  $\rightarrow$  Stop
a  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  inc  $\rightarrow$  c  $\rightarrow$  pass  $\rightarrow$  b  $\rightarrow$  fail  $\rightarrow$  Stop
...
b  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  pass  $\rightarrow$  a  $\rightarrow$  fail  $\rightarrow$  Stop
b  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  pass  $\rightarrow$  b  $\rightarrow$  fail  $\rightarrow$  Stop
b  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  pass  $\rightarrow$  c  $\rightarrow$  fail  $\rightarrow$  Stop
b  $\rightarrow$  inc  $\rightarrow$  e  $\rightarrow$  inc  $\rightarrow$  d  $\rightarrow$  pass  $\rightarrow$  a  $\rightarrow$  fail  $\rightarrow$  Stop
...
```

We define below a test driver that runs all these tests until a `pass` verdict, possibly followed by a `fail` verdict

is observed. This is ensured by a test monitor process $TMonitor$ executed alongside the tests.

$$Driver(SP, SUT) \downarrow_e = \left(\left\| \left\| T : Exhaust_T(SP) \downarrow_e \bullet Execution_{SUT}^{SP}(T) \right\| \right\| \{ \{ inc, pass, fail \} \} TMonitor \right)$$

where

$$TMonitor = inc \rightarrow TMonitor \square pass \rightarrow fail \rightarrow \mathbf{Stop}$$

Running $Driver(SP, SUT) \downarrow_e$ ensures that if e is executable by the SUT, it is covered at runtime. This means that some test based on a trace where e occurs reaches a conclusive verdict: *pass* or *fail*. (If the SUT is nondeterministic, the driver needs to be executed several times, and we need to rely on the complete test assumption.) The tests are executed in interleaving (in parallel independently) until one of them reaches a *pass* event. In this case, we either have an immediate deadlock, or just the associated forbidden continuation as an internal event of the test execution, followed by an observable *fail*. We recall that all events of the specification are internal to a test execution, and so not visible to the $TMonitor$ process.

A shortcoming of the above driver is that it makes an implicit arbitrary selection of a forbidden continuation. A more liberal $TMonitor$ should allow all (or specific) forbidden continuations to be tested.

If e is not covered at runtime with the strategy above, all the tests return an inconclusive verdict. On the other hand, the fact that all tests return an inconclusive verdict does not mean that e has not been executed since some of the tests may return an inconclusive verdict after having exercised e . \square

Selection of tests is often based on a given property π of traces or tests. Examples of such properties are coverage of an event, as considered above, or absence of a particular subsequence of events, and so on. In this case, it is usual to consider as selection criteria either: at least one test satisfying π is executed, or all the tests satisfying π are executed. We call the former existential selection of tests with respect to π , and the latter universal selection with respect to π . We can imagine, of course, some intermediate criteria where several of these tests are selected, or some different criterion that corresponds to a global property of the set of selected tests. As explained later on, however, these selection criteria can be addressed using compositions and variations of the definitions we describe below.

Properties of traces, rather than tests, may also be considered for selection criteria. In what follows, we first discuss existential or universal selection of traces and, in this case, the fact that there are several tests for one trace must be taken into account. Afterwards, we discuss selection based on properties of tests, and consider existential and universal selection as informally discussed above. We observe that the problem discussed in the previous section, which arises from non-implemented traces of the specification, is only an issue for existential selection (of either traces or tests) and its variants.

To characterise the criteria based on properties of traces, we note that for each trace t of a specification SP , we have the set of tests $AllTests_T(t, SP)$ defined below.

$$AllTests_T(t, SP) = \{ T_T(t, e) \mid e \notin initials(SP, t) \}$$

It contains all tests built from t and one of its forbidden continuations e .

For a property π of traces, we have the possible selection criteria below.

existential selection of traces: there is *at least one* trace t satisfying π , such that, $AllTests_T(t, SP)$, the test set identified by t is executed (that is, all tests in $AllTests_T(t, SP)$ are executed); or

universal selection of traces: for *all implemented* traces t satisfying π , $AllTests_T(t, SP)$ is executed.

Moreover, several properties π_1, \dots, π_n of traces may be considered; the coverage of a set of elements is an example. In such cases, one may require either of the following criteria.

multiple existential selection of traces: for every π_i , there is at least one trace t satisfying π_i , such that $AllTests_T(t, SP)$ is actually executed; or

multiple universal selection of traces: for every π_i , for every implemented trace t satisfying π_i , the set $AllTests_T(t, SP)$ is executed.

Coming back to test selection, for a given property π of tests, one may require either of the following.

existential selection of tests: there is *at least one* test T satisfying π , that is actually executed, if there is one that is implemented, or

universal selection of tests: *all implemented* tests T satisfying π are executed.

Direct test selection makes it possible to focus on specific forbidden continuations, for instance to consider only some specific forbidden event. The property π may also combine requirements on the traces and the forbidden continuations, like in, for instance, “all the tests with a trace t where event $e1$ occurs and event $e2$ is the forbidden continuation”. In Example 7, for instance, such a criterion for test selection allows us to restrict the set of relevant tests to those that have \mathbf{a} , for example, as a forbidden continuation. In this case, only three of the tests enumerated in Example 7 are retained for testing. Multiple existential and universal selection criteria for tests can be tackled in the same way as those for traces. Below, we consider exhaustive test sets and drivers for each of the above classes of traces and tests selection criteria.

3.2.1. Existential selection of traces

For existential selection, the selected test set $Existential_T^{trace}(SP) \upharpoonright_\pi$ needs to be structured as a set of test sets characterised by $AllTests_T$, as defined below.

$$Existential_T^{trace}(SP) \upharpoonright_\pi = \{AllTests_T(t, SP) \mid t \in traces(SP) \wedge t \vdash \pi\} \quad (1)$$

We use the notation $t \vdash \pi$ to indicate that the trace t satisfies the property π .

The execution of all the tests associated with an implemented trace t is specified below as a simple driver that runs independently and to conclusion all tests in $AllTests_T(t, SP)$.

$$\left\| \left\| T : AllTests_T(t, SP) \bullet Execution_{SUT}^{SP}(T) \right. \right.$$

A driver that ensures existential coverage can be specified as follows: independently, for each trace t that defines a set of tests $AllTests_T(t, SP)$ in $Existential_S(SP) \upharpoonright_\pi$, the driver above is run under the control of a monitor that observes whether a *pass* event takes place. When it does, it means that t is implemented. The monitor continues the execution of the other tests based on t , and sends to the other drivers for the other traces a *done* event to make them to stop. Formally, we have the following definition for a generic driver, for a specification SP and some property π of a trace.

$$\begin{aligned} ExistDriver^{trace}(SP, SUT) \upharpoonright_\pi = & \\ & \left(\left\| \left\| done \right\| \right\| TS : Existential_T^{trace}(SP) \upharpoonright_\pi \bullet \right. \\ & \left. \left(\left\| \left\| T : TS \bullet Execution_{SUT}^{SP}(T) \right\| \right\| \left\| inc, pass, fail \right\| \right\| TSMonitor \right) \setminus \left\| done \right\| \end{aligned}$$

where $TSMonitor$ is below. It stops when a *done* occurs, and raises a *done* event when a *pass* occurs.

$$TSMonitor = done \rightarrow Stop \square inc \rightarrow TSMonitor \square pass \rightarrow done \rightarrow Run$$

Run just continues running the other tests based on the same trace.

$$Run = inc \rightarrow Run \square pass \rightarrow Run \square fail \rightarrow Run$$

In the definition above of $ExistDriver^{trace}(SP) \upharpoonright_\pi$, each of the sets TS in $Existential_T^{trace}(SP) \upharpoonright_\pi$ is considered. The drivers for each of these sets are run in parallel, synchronising on the event *done*. Each driver runs all tests T in TS independently, but under the control of $TSMonitor$, which observes *done* as well as the verdict events. If a *done* event occurs, $TSMonitor$ stops (and so the tests under its control stop), because this indicates that a monitor for another test set has observed a test with a conclusive verdict. Otherwise, $TSMonitor$ and its tests proceed until a *pass* is observed, when *done* is raised, but the controlled tests proceed to conclusion; Run just ignores all verdict events. The *done* event is local to the existential driver; it is used only for synchronisation between the monitors of the tests associated with each trace t . Termination of the driver depends of the size of the set of selected tests. This is further discussed in Section 3.5.

Example 8. Coming back to coverage of an event e considered in Example 7, we have:

$$Existential_T^{trace}(SP) \upharpoonright_{(t \upharpoonright_e \neq \langle \rangle)} = \{ AllTests_T(t, SP) \mid t \in traces(SP) \wedge t \upharpoonright_e \neq \langle \rangle \}$$

The distributed union of the test sets in this set is exactly $Exhaust_T(SP) \upharpoonright_e$. Here, their grouping based on the traces t as defined by $AllTests_T(t, SP)$ allows us to consider the alternative driver below.

$$\begin{aligned} ExistDriver^{trace}(SP, SUT) \upharpoonright_{(t \upharpoonright_e \neq \langle \rangle)} = \\ \left(\left[\{ done \} \right] TS : Existential_T^{trace}(SP) \upharpoonright_{(t \upharpoonright_e \neq \langle \rangle)} \bullet \right. \\ \left. \left(\left[\left[T : TS \bullet Execution_{SUT}^{SP}(T) \right] \left[\{ inc, pass, fail \} \right] TSMonitor \right] \setminus \{ done \} \right) \right) \end{aligned}$$

This driver ensures that for at least one implemented trace covering e , if there is one, all the tests are executed, that is, all the forbidden events are attempted. As already said, the driver $Driver(SP, SUT) \upharpoonright_e$ presented in Example 7 attempts only one of them chosen arbitrarily. \square

As usual, if the SUT is nondeterministic, coverage has to rely on the complete test assumption and on running the driver several times. This is the case for all drivers that we present.

3.2.2. Universal selection of traces

The case of universal selection is easier to formalise; there is no need to consider a set of test sets as in the previous section. The test set to be considered is as follows.

$$Universal_T^{trace}(SP) \upharpoonright_\pi = \{ T_T(t, a) \mid t \in traces(SP) \wedge t \vdash \pi \wedge a \notin initials(SP, t) \} \quad (2)$$

All tests that are implemented must be run, so the driver attempts to execute all tests.

$$UnivDriver^{trace}(SP, SUT) \upharpoonright_\pi = \left[\left[T : Universal_T^{trace}(SP) \upharpoonright_\pi \bullet Execution_{SUT}^{SP}(T) \right] \right]$$

Example 9. Coming back again to coverage of an event e , we have

$$Exhaust_T(SP) \upharpoonright_e = Universal_T^{trace}(SP) \upharpoonright_{(t \upharpoonright_e \neq \langle \rangle)}$$

The corresponding driver $UnivDriver^{trace}(SP, SUT) \upharpoonright_{(s \upharpoonright_e \neq \langle \rangle)}$ executes all tests arising from all traces that include at least one occurrence of e . The driver $Driver(SP, SUT)$ in Example 7 executes only one of these tests for each trace with at least one occurrence of e . \square

3.2.3. Multiple existential or universal selection of traces

Selection criteria based on a set of properties $\Pi = \{\pi_1, \dots, \pi_n\}$ can be handled by considering the test sets and the drivers corresponding to each property π_i . For multiple existential coverage of traces, we have a set of test sets for each π_i , and the driver executes independently each of the existential drivers.

$$MExistDriver^{trace}(SP, SUT) \upharpoonright_\Pi = \left[\left[\pi_i : \Pi \bullet ExistDriver^{trace}(SP, SUT) \upharpoonright_{\pi_i} \right] \right]$$

Example 10. A classical example of a multiple existential selection criteria is that every event must be covered at least once. For a set E of events, we have the following driver:

$$MExistDriver^{trace}(SP, SUT) \upharpoonright_E = \left[\left[e : E \bullet ExistDriver^{trace}(SP, SUT) \upharpoonright_e \right] \right]$$

\square

For multiple universal selection criteria, we can consider a test set containing all tests obtained for each property π_i . There is no need to keep separate sets of tests, since they are all to be executed. We can, therefore, use a driver similar to $UnivDriver^{trace}(SP, SUT) \upharpoonright_\pi$ previously presented.

3.2.4. Existential or universal selection of tests

For a given property π of tests, we consider, for both existential and universal selection, the following exhaustive test set. We use $T \vdash \pi$ to denote the fact that the test T satisfies π .

$$Exhaust_T^{test}(SP) \upharpoonright_\pi = \{T : Exhaust_T(SP) \mid T \vdash \pi\}$$

There is no need to use a set of test sets like we have done for existential selection of traces.

For existential selection of the tests satisfying π , the driver is as follows.

$$ExistDriver^{test}(SP, SUT) \upharpoonright_\pi = \left(\left\| \left\| T : Exhaust_T^{test}(SP) \upharpoonright_\pi \bullet Execution_{SUT}^{SP}(T) \right\| \right\| \{inc, pass, fail\} \right) TMonitor$$

where $TMonitor$ is as in Example 7.

$$TMonitor = inc \rightarrow TMonitor \square pass \rightarrow fail \rightarrow \mathbf{Stop}$$

For universal selection of tests satisfying π , the driver is shown below.

$$UnivDriver^{test}(SP, SUT) \upharpoonright_\pi = \left\| \left\| T : Exhaust_T^{test}(SP) \upharpoonright_\pi \bullet Execution_{SUT}^{SP}(T) \right\| \right\|$$

This is exactly the same driver used for universal selection based on traces, except that it uses the tests in $Exhaust_T^{test}(SP) \upharpoonright_\pi$ (rather than those in $Exhaust_T^{trace}(SP) \upharpoonright_\pi$).

To give an example, we define below the process $end_T(T)$; it captures the final behaviour of a test T , where, after signalling a *pass*, it attempts to execute a forbidden continuation.

Definition 8.

$$end_T(pass \longrightarrow P) = P \qquad end_T(inc \longrightarrow e \longrightarrow P) = end_T(P)$$

This definition is used in the next example and in the next section.

Example 11. One may consider the universal selection criterion “all the tests where a given event e is a forbidden continuation”, possibly because it is especially critical not to execute this event e when it is forbidden by the specification. In this case, the test set is as follows.

$$Exhaust_T^{test}(SP) \upharpoonright_{\neg e} = \{T : Exhaust_T(SP) \mid end(T) = e \rightarrow fail \rightarrow \mathbf{Stop}\}$$

□

More imaginative selection criteria can be dealt with by the definition of adequate test sets or sets of test sets, or by composition of drivers. We have considered the main kinds of selection strategies, and shown how to construct the corresponding exhaustive test set for the required selection criteria and accompanying driver. We cannot claim to have explored every conceivable criterion, but note that driver composition with exchanges of messages, as done for existential selection, can cater for quite a wide range of selection strategies. If needed, it is possible to specify in *Circus* drivers that use, for example, counters and guards.

3.3. Factorisation

As already said, instead of using a monitor process to define special drivers, it is possible to factorise the tests in the exhaustive test set to define a single tree-shaped test as it is done for adaptive testing in [33]. This test offers the choice between all the tests selected; the second test in Example 6 illustrates the idea.

Factorisation decreases the number of inconclusive verdicts, since it gives the SUT the choice of several continuations available in the test set. Factorised tests have, however, a difficulty concerning coverage of behaviours, since the choices are no more under the control of the test and can be biased by the SUT.

Example 12. We consider, for instance, the factorisation of the following tests.

$$\begin{array}{l} \text{inc} \longrightarrow \text{a} \longrightarrow \text{pass} \longrightarrow \text{d} \longrightarrow \text{fail} \longrightarrow \text{Stop} \\ \text{inc} \longrightarrow \text{a} \longrightarrow \text{inc} \longrightarrow \text{e} \longrightarrow \text{pass} \longrightarrow \text{c} \longrightarrow \text{fail} \longrightarrow \text{Stop} \end{array}$$

It gives the following tree-shaped test.

$$\text{inc} \longrightarrow \text{a} \longrightarrow \left(\begin{array}{l} \text{pass} \longrightarrow \text{d} \longrightarrow \text{fail} \longrightarrow \text{Stop} \\ \square \\ \text{inc} \longrightarrow \text{e} \longrightarrow \text{pass} \longrightarrow \text{c} \longrightarrow \text{fail} \longrightarrow \text{Stop} \end{array} \right)$$

After the SUT performs the trace $\langle a \rangle$, there is a choice between the testing events *pass* and *inc* that is left up not to the test, but to the environment of the test execution. It is an external choice over events that are not in the alphabet of the SUT. \square

For the non-factorised tests, the environment of a test execution can be a simple process that is prepared to accept interaction on any of the events *inc*, *pass*, and *fail*, at any moment. Even in such a liberal environment, it is guaranteed that the interaction with the SUT defined by the trace that corresponds to the test is attempted in the test execution. If the test is factorised, though, the environment of the test execution is offered choices between *pass* and *inc* events, whenever there is a possibility, according to the specification, of extending the trace or deadlocking. To ensure coverage, both choices should be tried.

A related concern arises when there are several possibilities of failure.

Example 13. We consider, for instance, the factorisation of the following tests.

$$\begin{array}{l} \text{inc} \longrightarrow \text{a} \longrightarrow \text{pass} \longrightarrow \text{d} \longrightarrow \text{fail} \longrightarrow \text{Stop} \\ \text{inc} \longrightarrow \text{a} \longrightarrow \text{pass} \longrightarrow \text{e} \longrightarrow \text{fail} \longrightarrow \text{Stop} \end{array}$$

It gives the following tree-shaped test.

$$\text{inc} \longrightarrow \text{a} \longrightarrow \text{pass} \longrightarrow \left(\begin{array}{l} \text{d} \longrightarrow \text{fail} \longrightarrow \text{Stop} \\ \square \\ \text{e} \longrightarrow \text{fail} \longrightarrow \text{Stop} \end{array} \right)$$

In this case, after the occurrence of *a*, if the environment of the test execution chooses to synchronise on *pass*, it then has no control over the choice between *d* and *e*: these events are not in its alphabet. The test is not making a choice, but rather offering it to the SUT. If the SUT can perform any of them, the mistake is going to be detected, but it is not possible to determine which of the mistaken events are possible, since the SUT is not forced to perform any of them. \square

It remains, however, that the SUT drives the test by making the (external) choices of model events offered in the test. Since just one path of the tree-shaped test is executed in a test experiment, we have an existential selection (by the SUT) of a test combined into the tree. Moreover, only the complete test assumption and enough executions of the tree-shaped test can guarantee that all paths in the tree-shaped test (and, therefore, all tests combined in the tree) are executed.

We define below a function $Fact^{SP}(TS)$ that defines the factorised tree-shaped tests coming from a test set TS . It is defined in terms of a function $FactP^{SP}(TS, t)$, which takes as an extra parameter the path (a trace) t that defines the branch of the tree-shaped test that is being constructed.

Definition 9 (Factorisation).

$$Fact^{SP}(TS) = FactP^{SP}(TS, \langle \rangle)$$

$$FactP^{SP}(TS, t) = TFail(TS, t) \square TCont^{SP}(TS, t)$$

$TFail(TS, t) = \mathbf{Stop}$ [if $t \notin \text{trace}(\downarrow TS)$]

$TFail(TS, t) = \text{pass} \longrightarrow \square T : \{T : TS \mid \text{trace}_T(T) = t\} \bullet \text{end}_T(T)$ [if $t \in \text{trace}(\downarrow TS)$]

$TCont^{SP}(TS, t) = \mathbf{Stop}$ [if $\text{initials}_T(TS, t) = \emptyset$]

$TCont^{SP}(TS, t) = \text{inc} \longrightarrow \square e : \text{initials}_T(TS, t) \bullet e \longrightarrow \text{FactP}^{SP}(TS, t \hat{\ } \langle e \rangle)$ [if $\text{initials}_T(TS, t) \neq \emptyset$]

For any set TS , $\text{Fact}^{SP}(TS)$ is given by $\text{FactP}^{SP}(TS, \langle \rangle)$; initially, just the empty path has been constructed. For any path t , this is an external choice between tests defined by functions $TFail(TS, t)$ and $TCont^{SP}(TS, t)$. With $TFail(TS, t)$, we consider whether t is the trace of a test in TS or not. We define the set $\text{trace}(\downarrow TS)$ of traces of a set of tests TS as the set of specification traces used in its tests, not including the forbidden continuations. We use the relational image operator $f(\downarrow S)$; it provides the set of results obtained by applying f to each element in the set S . The trace of a single test T is defined as follows.

Definition 10.

$$\text{trace}_T(\text{pass} \longrightarrow P) = \langle \rangle \quad \text{trace}_T(\text{inc} \longrightarrow e \longrightarrow P) = \langle e \rangle \hat{\ } \text{trace}_T(P)$$

If t is not a trace of a test in TS , then $TFail(TS, t)$ is **Stop**, which is a unit for external choice: $\mathbf{Stop} \square P = P$. For example, TS may not have any test for the empty trace $\langle \rangle$, and in this case $TFail(TS, \langle \rangle)$ does not add anything to the tree-shaped test. On the other hand, if there are tests T in TS whose trace is t , then, for each of them, we add its treatment of the forbidden continuation, characterised by $\text{end}_T(T)$, to the tree.

With $TCont^{SP}(TS, t)$, we consider whether t is the prefix of a trace of a test in TS or not. We use the notion of initials of a test set TS after a trace t formally defined as follows.

Definition 11.

$$\text{initials}_T(TS, t) = \{e \mid \exists T : TS; s \bullet \text{trace}_T(T) = t \hat{\ } \langle e \rangle \hat{\ } s\}$$

If t has no continuation, then $TCont^{SP}(TS, t)$ is just **Stop**. Otherwise, we add an intermediary verdict *inc* followed by a branch (in an external choice) for each of the continuations e .

Example 14. *The test in Example 6 can be built from the test set including*

$\text{inc} \longrightarrow \mathbf{a} \longrightarrow \text{inc} \longrightarrow \mathbf{e} \longrightarrow \text{pass} \longrightarrow \mathbf{d} \longrightarrow \text{fail} \longrightarrow \mathbf{Stop}$

and

$\text{inc} \longrightarrow \mathbf{b} \longrightarrow \text{inc} \longrightarrow \mathbf{e} \longrightarrow \text{pass} \longrightarrow \mathbf{c} \longrightarrow \text{fail} \longrightarrow \mathbf{Stop}$

No special driver is required in this case. □

As already indicated above, however, the two approaches discussed here, namely, factorisation and the definition of exhaustive test sets and drivers, are not equivalent. In the execution of a factorised test, the control, exercised via choice of events, is left to the SUT or the test environment. This is in contrast with the drivers in Section 3.2, which control the execution of the tests to which they are associated.

A third approach is to alternate the control between the driver and the SUT, as in on-the-fly testing [58, 59, 32], where test derivation and execution are integrated into a single process. This approach is also called online testing, in contrast with offline testing, where test generation is a separate process that takes place before execution. Online testing is especially relevant when testing reactive systems: the next event of the tester is generated in function of the reaction of the SUT.

In this paper, we have not addressed the issue of generation of the various exhaustive test sets occurring in the drivers: they are defined in an abstract, predicative way. Similarly, the factorisation function above starts from a given subset of the exhaustive test set, and thus is not a generation algorithm. The corresponding

test-generation algorithms from a *Circus* specification can be developed and used either offline, following the patterns of the drivers defined in Section 3.2, or online.

Online testing would avoid running in parallel all the tests that satisfy the criterion. On the other hand, online testing would require integrating in the driver the generation algorithm of the exhaustive test sets for ensuring that all the relevant tests are attempted until one or all are completed (depending of the kind of considered selection: existential or universal).

3.4. Partial implementation of symbolic tests

As explained in Section 2.3, the *Circus* testing theory offers a symbolic characterisation of the concrete tests and of the exhaustive concrete test set. It is possible to specify a testing strategy for a *Circus* specification at the concrete level, just following what is presented in Section 3.2 or 3.3. Uniformity hypotheses on instantiations of constrained symbolic tests, however, provide a natural selection criterion. For instance, a simple hypothesis states that the verdict for all instances of a symbolic test is the same. In this case, we assume that the verdict of the test in Example 2 is the same verdict of all instances of the symbolic test in Example 5. Constrained symbolic tests enable the definition of powerful selection strategies that take into account data types properties. We, therefore, consider now how to select from symbolic tests, and then instantiate the selected tests and run them via concrete drivers to ensure runtime coverage.

We have a two-level selection method. Firstly, we need to provide a symbolic exhaustive test set with respect to the selection criteria of interest. Just like in Section 3.2, this amounts to the definition of all those symbolic tests that, when instantiated, may contribute at runtime to the satisfaction of the criterion. In this case, we need to rely on properties of symbolic traces or symbolic tests, but that is all.

Secondly, we need to define drivers for the instantiations of the symbolic tests. The challenge is that, when starting the testing experiments with a given symbolic test, it may be only partially implemented. Moreover, if there is no implemented instantiation, it is necessary to try another symbolic test of the exhaustive test set, and actually all of them until either one concrete test reaches a conclusion, or it is certain that it is impossible to fulfill the criterion for the SUT. Mechanisms similar to those in Section 3.2 must be used for existential or universal selection among those concrete tests corresponding to a symbolic test.

The execution of all the implemented instances of a given symbolic test ST is specified as follows.

$$UnivInstDriver(ST, SP, SUT) = \left\| \left\| T : instTest(ST) \bullet Execution_{SUT}^{SP}(T) \right. \right.$$

As said above, however, symbolic tests are generally used to select only one of their instantiations. The execution of at least one implemented instance of a symbolic test ST , if there is one, is specified below.

$$ExistInstDriver(ST, SP, SUT) = \left(\left\| \left\| T : instTest(ST) \bullet Execution_{SUT}^{SP}(T) \right. \right. \llbracket \{ inc, pass, fail \} \rrbracket TMonitor$$

This is similar to the existential selection driver addressed in Section 3.2, but we monitor tests coming from a single test set, namely, the set of instantiations of the symbolic test ST , rather than a set of tests. As soon as a test T in the set $instTest(ST)$ reaches a *pass* verdict, the monitor stops all other tests. The test that reached a *pass* is allowed to conclude (by possibly giving a later *fail* verdict) before stopping.

The two specifications above can be realised by the use at runtime of some adequate solver for enumerating instantiations in a way similar to [56] and [3]. In this case, the instances are generated one at a time, until a *pass* event is observed; only the instances of the same symbolic test that generated the current concrete test are then considered, and no other symbolic tests are tried.

In summary, in a selection strategy for symbolic tests, we can have existential or universal selection of symbolic traces or tests, with an existential or universal selection of instances of symbolic tests. In combining the results in Section 3.2 with $UnivInstDriver(ST, SP, SUT)$ or $ExistInstDriver(ST, SP, SUT)$ presented above, what we need to do is to use these drivers as the characterisation of the execution of a symbolic test. For example, if we have an existential selection of symbolic traces (based on a property π), the set

$Existential_T^{trace}(SP) \downarrow_\pi$ as defined in Section 3.2 is a set of sets of symbolic tests. In this case, if universal selection of instances is required, the existential driver to be used should be as follows.

$$ExistUnivDriver^{cstrace}(SP, SUT) \downarrow_\pi = \left(\begin{array}{l} \llbracket \{ \text{done} \} \rrbracket TS : Existential_T^{trace}(SP) \downarrow_\pi \bullet \\ \left(\left\| \left\| T : TS \bullet UnivInstDriver(T, SP, SUT) \right\| \llbracket \{ \text{inc}, \text{pass}, \text{fail} \} \rrbracket TSMonitor \right\| \right) \setminus \{ \text{done} \} \end{array} \right)$$

This is the existential driver from Section 3.2, but the process that captures execution of a test is not $Execution_{SUT}^{SP}(T)$, but the universal driver $UnivInstDriver(T, SP, SUT)$ for symbolic tests defined above. Each TS is a set of symbolic tests for the same symbolic trace. Each T in TS is, therefore, a symbolic test, which is executed by $UnivInstDriver(T, SP, SUT)$. As soon as an instance of a symbolic test reaches a *pass*, just as for concrete tests, the $TSMonitor$ stops the tests for (the instances of) symbolic tests in the other sets of symbolic tests. On the other hand, all instances of all other symbolic tests in the same set, which cover all symbolic forbidden continuations, are allowed to proceed to completion.

If, on the other hand, existential selection of instances is required, the driver to be used is as follows.

$$ExistExistDriver^{cstrace}(SP, SUT) \downarrow_\pi = \left(\begin{array}{l} \llbracket \{ \text{done} \} \rrbracket TS : Existential_T^{trace}(SP) \downarrow_\pi \bullet \\ \left(\left\| \left\| T : TS \bullet ExistInstDriver(T, SP, SUT) \right\| \llbracket \{ \text{inc}, \text{pass}, \text{fail} \} \rrbracket TSMonitor \right\| \right) \setminus \{ \text{done} \} \end{array} \right)$$

In this case, as soon as an instance of a symbolic test reaches a *pass*, the $TSMonitor$ stops the tests for the other sets of symbolic tests. On the other hand, one instance of all other symbolic tests in the same set is allowed to complete. So, for each forbidden continuation, at least one instance is executed.

Example 15. For illustration, we consider $Existential_T^{trace}(SP) \downarrow_\pi$ to be a set containing just two sets of symbolic tests. The first set contains tests for the symbolic trace $(\langle \mathbf{a}.\alpha, \alpha \in \{6, 7\} \rangle)$; we assume that the forbidden continuations for this trace are $(\mathbf{a}.\beta, \alpha \in \{6, 7\} \wedge \alpha = \beta)$ and $(\mathbf{b}, \alpha \in \{6, 7\})$. The second set of tests is for the empty trace $(\langle \rangle, \mathbf{true})$ with the single forbidden continuation $(\mathbf{out}.\alpha, \alpha = 5)$. The first set has two tests, one for each forbidden continuation, and the second has just one symbolic test.

$$Existential_T^{trace}(SP) \downarrow_\pi = \left\{ \begin{array}{l} \{ \text{inc} \rightarrow \mathbf{a}.\alpha : \alpha \in \{6, 7\} \rightarrow \text{pass} \rightarrow \mathbf{a}.\beta : \alpha \in \{6, 7\} \wedge \alpha = \beta \rightarrow \text{fail} \rightarrow \mathbf{Stop}, \\ \text{inc} \rightarrow \mathbf{a}.\alpha : \alpha \in \{6, 7\} \rightarrow \text{pass} \rightarrow \mathbf{b} : \alpha \in \{6, 7\} \rightarrow \text{fail} \rightarrow \mathbf{Stop} \}, \\ \{ \text{pass} \rightarrow \mathbf{out}.\alpha : \alpha = 5 \rightarrow \text{fail} \rightarrow \mathbf{Stop} \} \end{array} \right\}$$

Each test in the first set has two instantiations: one where α (and β) takes the value 6, and another where it takes the value 7. There is only one instantiation of the test in the second set. The driver for universal instantiation is the following process as defined by $ExistUnivDriver^{cstrace}(SP, SUT) \downarrow_\pi$.

$$\left(\left(\left(\begin{array}{l} Execution_{SUT}^{SP}(\text{inc} \rightarrow a.6 \rightarrow \text{pass} \rightarrow a.6 \rightarrow \text{fail} \rightarrow \mathbf{Stop}) \\ \parallel \\ Execution_{SUT}^{SP}(\text{inc} \rightarrow a.7 \rightarrow \text{pass} \rightarrow a.7 \rightarrow \text{fail} \rightarrow \mathbf{Stop}) \\ \parallel \\ Execution_{SUT}^{SP}(\text{inc} \rightarrow a.6 \rightarrow \text{pass} \rightarrow b \rightarrow \text{fail} \rightarrow \mathbf{Stop}) \\ \parallel \\ Execution_{SUT}^{SP}(\text{inc} \rightarrow a.7 \rightarrow \text{pass} \rightarrow b \rightarrow \text{fail} \rightarrow \mathbf{Stop}) \\ \llbracket \{ \text{inc}, \text{pass}, \text{fail} \} \rrbracket \\ TSMonitor \\ \llbracket \{ \text{done} \} \rrbracket \\ \left(\begin{array}{l} Execution_{SUT}^{SP}(\text{pass} \rightarrow \text{out}.5 \rightarrow \text{fail} \rightarrow \mathbf{Stop}) \\ \llbracket \{ \text{inc}, \text{pass}, \text{fail} \} \rrbracket \\ TSMonitor \end{array} \right) \end{array} \right) \right) \setminus \{ \text{done} \}$$

In this case, when one $TSMonitor$ detects a *pass* event, it stops the other monitor and its tests, but lets all

its other tests run to conclusion. The driver for existential instantiation, on the other hand, is as follows.

$$\left(\left(\left(\left(\left(\begin{array}{c} Execution_{SUT}^{SP}(inc \longrightarrow a.6 \longrightarrow pass \longrightarrow a.6 \longrightarrow fail \longrightarrow \mathbf{Stop}) \\ \parallel \\ Execution_{SUT}^{SP}(inc \longrightarrow a.7 \longrightarrow pass \longrightarrow a.7 \longrightarrow fail \longrightarrow \mathbf{Stop}) \\ \llbracket \{ inc, pass, fail \} \rrbracket \\ TMonitor \end{array} \right) \right) \right) \right) \right) \left(\left(\left(\left(\begin{array}{c} Execution_{SUT}^{SP}(inc \longrightarrow a.6 \longrightarrow pass \longrightarrow b \longrightarrow fail \longrightarrow \mathbf{Stop}) \\ \parallel \\ Execution_{SUT}^{SP}(inc \longrightarrow a.7 \longrightarrow pass \longrightarrow b \longrightarrow fail \longrightarrow \mathbf{Stop}) \\ \llbracket \{ inc, pass, fail \} \rrbracket \\ TMonitor \end{array} \right) \right) \right) \right) \setminus \{ done \} \\ \left(\begin{array}{c} \llbracket \{ inc, pass, fail \} \rrbracket \\ TSMonitor \\ \llbracket \{ done \} \rrbracket \\ Execution_{SUT}^{SP}(pass \longrightarrow out.5 \longrightarrow fail \longrightarrow \mathbf{Stop}) \\ \llbracket \{ inc, pass, fail \} \rrbracket \\ TSMonitor \end{array} \right)$$

In this case, once a *TSMonitor* detects a *pass*, it stops the other *TSMonitor* processes as before, and the *TMonitor* processes under its control allow one test of each set of instantiations to conclude. \square

For universal selection of symbolic tests, again, where $Execution_{SUT}^{SP}(T)$ is used in the driver in Section 3.2, we can use either $UnivInstDriver(T, SP, SUT)$ or $ExistInstDriver(T, SP, SUT)$. Accordingly, we either have all instances or one instance of all symbolic tests run.

If the execution of one instance of one symbolic test is required, what we need is existential selection of symbolic tests, rather than symbolic traces. For the construction of the driver, we proceed as shown above.

We observe that the test drivers presented above can also handle the variant notion of instantiation with symbolic final events defining synchronisations with predicates mentioned at the end of Section 2.3.

3.5. Finiteness

Selection criteria must identify a finite number of tests for execution. The exhaustive test sets for a selection criteria, however, are not necessarily finite.

We have suggested the use of classical uniformity and regularity hypotheses [4]. Uniformity hypotheses are based on the definition of a finite collection of test subsets (uniformity domains), with the assumption that in each subset the tests succeed or fail uniformly. Although there is a finite number of uniformity domains, each may have an infinite number of (symbolic) tests that can contribute to coverage of that domain. In fact, uniformity hypotheses do not eliminate tests from the exhaustive test set; they just group them to form the uniformity domains since they are multiple existential selection criterion (see Section 3.2.3).

Regularity hypotheses are based on a bound of some metrics on the tests: their length, number of occurrences of some events, or number of couples of events, for instance. Typically, a regularity hypothesis eliminates tests from the exhaustive test set. In the presence of an infinite number of events, like in typical *Circus* specifications, regularity hypotheses may not be enough to ensure finiteness. For instance, with an infinite number of events it is possible to construct an infinite number of traces of a given size.

The need for drivers that potentially try all tests in an exhaustive test set, as explained above, can lead to drivers that do not necessarily terminate if the exhaustive test set is infinite. In the next section, we give an example of a sophisticated selection criterion for symbolic tests. We use it in Section 4.5 in the context of our framework, and show how we select finite test sets for that example.

4. Synchronisation coverage in *Circus*

So far, all our examples are based on generic properties of traces or tests for selection. In this section, we consider a more specific selection criterion, based on the coverage of synchronisations between parallel

processes of a *Circus* model. To express this criterion formally, we need to identify traces of the model that might lead to at least one such synchronisation. This is the non-trivial task that we address in this section. We identify the relevant traces of the specification from traces of its parallel components. In addition, we prove that the set that we characterise is indeed a set of traces of the specification and, therefore, adequate for use in the context of our framework to define exhaustive test sets. Our formalisation and proof is in the context of the symbolic operational semantics of *Circus*, and defines how to deal with constraints and alphabets of symbolic variables in constrained symbolic traces.

The need for specific notions of coverage for concurrent programs was recognised very early. The first precise propositions were published at the end of the eighties, when the first general concurrent programming languages appeared (see Section 5 and [55]). More specifically, synchronisation coverage was discussed in [31], where a coverage criterion based on the execution of sequences of communications is proposed, and in [5] where synchronisation coverage for Ada programs is defined and instrumented. More recent work [48, 47] has focussed on integration testing based on coverage of communications between components of distributed real-time systems specified using deterministic timed interface automata. We take inspiration from these works and consider coverage of communications (synchronisations) between parallel processes in *Circus*.

In Section 4.1, we define the set $ScstraceP_{cs}^{(a_1, a_2)}(P_1, P_2)$ of pairs of constrained symbolic traces of parallel processes P_1 and P_2 that can lead to their synchronisation on channels in the set cs . We use this to define the set $Scstraces_{cs}^{a_1, a_2}(P_1, P_2)$ of synchronisation traces in Section 4.2. In Section 4.3, we generalise our definitions for parallelisms involving more than two processes. Finally, in Section 4.4, we define the set $selectS^a(P)$ of synchronisation traces of any *Circus* model P . Our main result in this section establishes that this is a subset of the traces of P , as required for the selection strategies discussed in the previous section.

4.1. Synchronisation pairs

We consider parallel processes P_1 and P_2 and their sets of constrained symbolic traces $cstraces^{a_1}(P_1)$ and $cstraces^{a_2}(P_2)$. We identify below the sets of pairs of traces from each process that can lead to their synchronisation, at least once, when they are executed in parallel with a synchronisation set cs .

Definition 12 (Synchronising constrained symbolic trace pairs). *For processes P_1 and P_2 , synchronisation channel set cs , and disjoint alphabets a_1 and a_2 , we define $ScstraceP_{cs}^{(a_1, a_2)}(P_1, P_2)$ as follows.*

$$ScstraceP_{cs}^{(a_1, a_2)}(P_1, P_2) = \left\{ \begin{array}{l} \mathbf{cst}_1 : cstraces^{a_1}(P_1); \mathbf{cst}_2 : cstraces^{a_2}(P_2) \mid \\ (\mathbf{cst}_1 \upharpoonright_C cs) \neq \langle \rangle \wedge (\mathbf{cst}_2 \upharpoonright_C cs) \neq \langle \rangle \wedge \text{satisfiable}_{cs}(\mathbf{cst}_1, \mathbf{cst}_2) \\ \end{array} \right\}$$

A synchronising pair of traces \mathbf{cst}_1 and \mathbf{cst}_2 is characterised in terms of their projections to communications over channels in cs and their joint satisfiability $\text{satisfiable}_{cs}((\mathbf{st}_1, \mathbf{c}_1), (\mathbf{st}_2, \mathbf{c}_2))$. To avoid naming conflicts in the traces of the pair, we consider disjoint alphabets a_1 and a_2 : those with no common elements.

Example 16. *We list some of the constrained symbolic traces of the process *Components* in Section 2.1.*

$$(\langle \text{in.}\alpha_1, \text{left.}\beta_1 \rangle, \alpha_1 = \beta_1) \tag{3}$$

$$(\langle \text{in.}\alpha_1, \text{left.}\beta_1, \text{right.}\gamma_1, \text{left.}\delta_1 \rangle, \alpha_1 = \beta_1 \wedge \beta_1 = \delta_1) \tag{4}$$

$$(\langle \text{in.}\alpha_1, \text{left.}\beta_1, \text{right.}\gamma_1, \text{left.}\delta_1, \text{right.}\epsilon_1 \rangle, \alpha_1 = \beta_1 \wedge \beta_1 = \delta_1) \tag{5}$$

*For the *Medium* process, some traces are as follows.*

$$(\langle \text{left.}\alpha_2 \rangle, \text{true}) \tag{6}$$

$$(\langle \text{left.}\alpha_2, \text{right.}\beta_2 \rangle, \alpha_2 = \beta_2) \tag{7}$$

$$(\langle \text{left.}\alpha_2, \text{right.}\beta_2, \text{left.}\gamma_2 \rangle, \alpha_2 = \beta_2) \tag{8}$$

$$(\langle \text{left.}\alpha_2, \text{right.}\beta_2, \text{left.}\gamma_2 \rangle, 1 - \alpha_2 = \beta_2) \tag{9}$$

$$(\langle \text{left.}\alpha_2, \text{right.}\beta_2, \text{left.}\gamma_2, \text{right.}\delta_2 \rangle, \alpha_2 = \beta_2 \wedge 1 - \delta_2 = \gamma_2) \tag{10}$$

$$(\langle \text{left.}\alpha_2, \text{right.}\beta_2, \text{left.}\gamma_2, \text{right.}\delta_2 \rangle, 1 - \alpha_2 = \beta_2 \wedge \delta_2 = \gamma_2) \tag{11}$$

*The synchronising pairs for the parallelism between *Components* and *Medium* cannot include the empty*

trace, or the traces that do not have at least one occurrence of an event representing a communication on left or right. Additionally, for example, the traces (3) and (6) above form a synchronising pair, but (3) and (7) do not, because it is not possible to match the event on right in (7): this pair of traces is not satisfiable. The traces (5) and (10), on the other hand, also form a synchronising pair.

In addition, for illustration, we consider a different action for the Receiver process in Section 2.1 used in the definition of Components. Namely, we consider $right?m1 \rightarrow right.m1 \rightarrow \dots$, instead of $right?m1 \rightarrow right?m2 \rightarrow \dots$. In this case, once the message $m1$ is input only another copy of $m1$ is accepted in the second communication via $right$: that is, instead of a new input $right?m2$, we have a synchronisation $right.m1$ based on $m1$. So, the constraint of the trace (5) above would be $\alpha_1 = \beta_1 \wedge \beta_1 = \delta_1 \wedge \epsilon_1 = \gamma_1$. Such a trace would not match the trace (10), because matching requires $\epsilon_1 = \delta_2$. This is not consistent with the requirement that $\epsilon_1 = \gamma_1$ because matching also requires $\gamma_1 = 1 - \delta_2$. \square

In Definition 12, we require that the projections of the traces are not empty, since otherwise the pair of traces does not really lead to any synchronisation. (In fact, it is enough to require that one of them is not empty, since satisfiability requires them to be of the same length.) The projection ($\mathbf{cst} \downarrow_{\mathbf{C}} \mathbf{cs}$) of a constrained symbolic trace \mathbf{cst} to the set \mathbf{cs} is the sequence of channels (rather than events) obtained by keeping from \mathbf{cst} only the channels from events representing communications through a channel in \mathbf{cs} .

Example 17. In our example, the projections of (3) and (6) to $\{\text{left}, \text{right}\}$, for example, are both $\langle \text{left} \rangle$. The projection of (4) to the same set of channels is $\langle \text{left}, \text{right}, \text{left} \rangle$. Finally, the projection of (10) is $\langle \text{left}, \text{right}, \text{left}, \text{right} \rangle$. \square

Formal definitions for this operator and others used later are in Appendix A.

Satisfiability for pairs of (synchronising) constrained symbolic traces embeds the fact that the constraints of the traces are jointly satisfiable, even when we match the communications on the synchronisations that they require. This matching gives rise to equalities that strengthen the requirements imposed by the original individual constraints, which must be satisfiable as well.

Definition 13 (Satisfiability). For every pair $(\mathbf{st}_1, \mathbf{c}_1)$ and $(\mathbf{st}_2, \mathbf{c}_2)$ of constrained symbolic traces and synchronisation channel set \mathbf{cs} , we define $\text{satisfiable}_{\mathbf{cs}}((\mathbf{st}_1, \mathbf{c}_1), (\mathbf{st}_2, \mathbf{c}_2))$ as follows.

$$\text{satisfiable}_{\mathbf{cs}}((\mathbf{st}_1, \mathbf{c}_1), (\mathbf{st}_2, \mathbf{c}_2)) \Leftrightarrow \exists \mathbf{a}_1, \mathbf{a}_2 \bullet (\mathbf{st}_1 \downarrow_{\mathbf{E}} \mathbf{cs}) =_{\mathbf{S}} (\mathbf{st}_2 \downarrow_{\mathbf{E}} \mathbf{cs}) \wedge \mathbf{c}_1 \wedge \mathbf{c}_2$$

where \mathbf{a}_1 and \mathbf{a}_2 are the alphabets of $(\mathbf{st}_1, \mathbf{c}_1)$ and $(\mathbf{st}_2, \mathbf{c}_2)$.

The equalities are defined by the operator $\mathbf{st}_3 =_{\mathbf{S}} \mathbf{st}_4$, which identifies the requirement for synchronisation of the pair of symbolic traces \mathbf{st}_3 and \mathbf{st}_4 containing only communications that need to synchronise. In the above definition, these traces are $\mathbf{st}_1 \downarrow_{\mathbf{E}} \mathbf{cs}$ and $\mathbf{st}_2 \downarrow_{\mathbf{E}} \mathbf{cs}$. The projection $\mathbf{st} \downarrow_{\mathbf{E}} \mathbf{cs}$ is similar to $\mathbf{cst} \downarrow_{\mathbf{C}} \mathbf{cs}$, but it applies to symbolic traces \mathbf{st} and results in symbolic traces: sequences of events, rather than of channels.

Example 18. The projection of the symbolic trace (4) in Example 16 to the set $\{\text{left}, \text{right}\}$, for example, is $\langle \text{left}.\beta_1, \text{right}.\gamma_1, \text{left}.\delta_1 \rangle$. For (10) the result of the projection is this same trace. \square

As formalised below, the result of $\mathbf{st}_1 =_{\mathbf{S}} \mathbf{st}_2$ is a conjunction of equalities that require that the values of the symbolic variables used to represent synchronising communications in \mathbf{st}_1 and \mathbf{st}_2 are the same.

Definition 14 (Conjunction of symbolic equalities). We define inductively $\mathbf{st}_1 =_{\mathbf{S}} \mathbf{st}_2$ as follows.

$$\begin{aligned} \langle \rangle &=_{\mathbf{S}} \langle \rangle = \mathbf{true} \\ \langle \rangle &=_{\mathbf{S}} \mathbf{st}_2 = \mathbf{false} \\ \mathbf{st}_1 &=_{\mathbf{S}} \langle \rangle = \mathbf{false} \\ \langle \langle \mathbf{d}.\alpha \rangle \wedge \mathbf{st}_1 \rangle &=_{\mathbf{S}} \langle \langle \mathbf{d}.\beta \rangle \wedge \mathbf{st}_2 \rangle = \alpha = \beta \wedge (\mathbf{st}_1 =_{\mathbf{S}} \mathbf{st}_2) \\ \langle \langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1 \rangle &=_{\mathbf{S}} \langle \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2 \rangle = \mathbf{false}, \text{ if } \mathbf{d}_1 \neq \mathbf{d}_2 \end{aligned}$$

If \mathbf{st}_1 and \mathbf{st}_2 have different lengths or require communications over different channels, synchronisation is not possible and the requirement defined by $\mathbf{st}_1 =_{\mathbf{S}} \mathbf{st}_2$ is just **false**.

Example 19. For the symbolic traces of (3) and (6), we get the equality $\beta_1 = \alpha_2$. For (3) and (7), we get just false. For (5) and (10), we get the conjunction of equalities $\beta_1 = \alpha_2 \wedge \gamma_1 = \beta_2 \wedge \delta_1 = \gamma_2 \wedge \epsilon_1 = \delta_2$. \square

As seen in Definition 13, in $\text{satisfiable}_{\text{cs}}((\text{st}_1, \text{c}_1), (\text{st}_2, \text{c}_2))$, the variables in the alphabets \mathbf{a}_1 and \mathbf{a}_2 are used in $(\text{st}_1 \upharpoonright_{\text{E}} \text{cs}) =_S (\text{st}_2 \upharpoonright_{\text{E}} \text{cs})$ and in the constraints c_1 and c_2 .

4.2. Synchronisation traces

Using the notion of synchronisation pairs, we define the notion of synchronisation traces, which are constrained symbolic traces that record a parallel execution of a pair of processes (rather than separate executions of the parallel processes captured by synchronisation pairs).

Example 20. Continuing with our example, from the synchronising pair of traces (3) and (6), we get the synchronising trace $(\langle \text{in}.\alpha, \text{left}.\beta \rangle, \alpha = \beta)$ over the alphabet $\langle \alpha, \beta, \gamma, \dots \rangle$ to capture possible observations of the parallel execution of *Components* and *Medium*. From (5) and (10), we get the synchronising trace $(\langle \text{in}.\alpha, \text{left}.\beta, \text{right}.\gamma, \text{left}.\delta, \text{right}.\epsilon \rangle, \alpha = \beta \wedge \beta = \gamma \wedge \beta = \delta \wedge \delta = 1 - \epsilon)$. \square

As defined below, synchronisation traces are characterised by the sets $\text{Scstraces}_{\text{cs}}^{\mathbf{a}, (\mathbf{a}_1, \mathbf{a}_2)}(\text{P}_1, \text{P}_2)$ of constrained symbolic traces, formed from the synchronising constrained symbolic trace pairs in $\text{ScstraceP}_{\text{cs}}^{\mathbf{a}_1, \mathbf{a}_2}(\text{P}_1, \text{P}_2)$, which, as formalised in Definition 12, are pairs of traces for the processes P_1 and P_2 , using disjoint alphabets \mathbf{a}_1 and \mathbf{a}_2 , and based on the synchronisation set cs .

Definition 15 (Synchronising traces). For every pair of processes P_1 and P_2 , synchronisation channel set cs , alphabets \mathbf{a}_1 , \mathbf{a}_2 , and \mathbf{a} , we define the set $\text{Scstraces}_{\text{cs}}^{\mathbf{a}, (\mathbf{a}_1, \mathbf{a}_2)}(\text{P}_1, \text{P}_2)$ over the alphabet \mathbf{a} as follows.

$$\text{Scstraces}_{\text{cs}}^{\mathbf{a}, (\mathbf{a}_1, \mathbf{a}_2)}(\text{P}_1, \text{P}_2) = \bigcup \{ \mathbf{p} : \text{ScstraceP}_{\text{cs}}^{\mathbf{a}_1, \mathbf{a}_2}(\text{P}_1, \text{P}_2) \bullet \mathcal{N}^{\mathbf{a}}(\mathbf{p}.1 \parallel \text{cs} \parallel \mathbf{p}.2) \}$$

To form traces out of the pairs of traces \mathbf{p} in $\text{ScstraceP}_{\text{cs}}^{\mathbf{a}_1, \mathbf{a}_2}(\text{P}_1, \text{P}_2)$, we apply the synchronisation operator $(\text{st}_1, \text{c}_1) \parallel \text{cs} \parallel (\text{st}_2, \text{c}_2)$ (Definition 16) to the traces $\mathbf{p}.1$ and $\mathbf{p}.2$ in \mathbf{p} . We use $t.i$ to refer to the i -th element of a tuple t . The synchronisation operator identifies all the traces that can arise from the parallel execution of processes that can carry out the traces $\mathbf{p}.1$ and $\mathbf{p}.2$, synchronising on the channels in cs .

Example 21. In the case of our example, from the synchronising pair of traces (3) and (6), we get the synchronising trace $(\langle \text{in}.\alpha_1, \text{left}.\beta_1 \rangle, \alpha_1 = \beta_1)$. From (5) and (10), we get a synchronising trace $(\langle \text{in}.\alpha_1, \text{left}.\alpha_2, \text{right}.\beta_2, \text{left}.\gamma_2, \text{right}.\delta_2 \rangle, \alpha_1 = \alpha_2 \wedge \alpha_2 = \beta_2 \wedge \alpha_2 = \gamma_2 \wedge \gamma_2 = 1 - \delta_2)$. \square

As illustrated above, the result of the application of $\mathbf{p}.1 \parallel \text{cs} \parallel \mathbf{p}.2$, as formalised in Definition 16, is a set of traces with possibly different alphabets. To obtain a set of traces over the given alphabet \mathbf{a} , we use in Definition 15 the normalising function $\mathcal{N}^{\mathbf{a}}$ (Definition 18), which basically carries out a renaming.

Synchronisation for constrained symbolic traces is specified in Definition 16 in terms of the similar operator $\text{st}_1 \parallel \text{cs}, \text{c} \parallel \text{st}_2$ introduced in Definition 17. This new operator is for symbolic traces st_1 and st_2 and has an extra parameter: a constraint c . In the specification of $(\text{st}_1, \text{c}_1) \parallel \text{cs} \parallel (\text{st}_2, \text{c}_2)$ in Definition 16, it is applied with parameter $\text{c}_1 \wedge \text{c}_2$ to capture the original constraints.

Example 22. From the synchronising pair of traces (3) and (6), the synchronisation operator for symbolic traces gives the constrained symbolic trace $(\langle \text{in}.\alpha_1, \text{left}.\beta_1 \rangle, \alpha_1 = \beta_1 \wedge \beta_1 = \alpha_2)$. Similarly, from (5) and (10), we get a trace with the symbolic trace $\langle \text{in}.\alpha_1, \text{left}.\beta_1, \text{right}.\gamma_1, \text{left}.\delta_1, \text{right}.\epsilon_1 \rangle$, and with a constraint that conjoins the original constraints of (5) and (10) and a conjunction of equalities $\beta_1 = \alpha_2 \wedge \gamma_1 = \beta_2 \wedge \delta_1 = \gamma_2 \wedge \epsilon_1 = \delta_2$ that forces the matching of their synchronised events. \square

As illustrated above, synchronisation of symbolic traces gives rise to constraints that refer to variables that are not part of the resulting symbolic trace. We, therefore, in defining $(\text{st}_1, \text{c}_1) \parallel \text{cs} \parallel (\text{st}_2, \text{c}_2)$ below, use the \mathcal{R} function to quantify away the spurious symbolic variables, which are used in the traces of the parallel processes, but are not needed in the synchronisation trace resulting from them.

Definition 16 (Synchronisation operator - constrained symbolic traces). We define the set of synchronising constrained symbolic traces $(\mathbf{st}_1, \mathbf{c}_1) \llbracket \mathbf{cs} \rrbracket (\mathbf{st}_2, \mathbf{c}_2)$ as follows.

$$(\mathbf{st}_1, \mathbf{c}_1) \llbracket \mathbf{cs} \rrbracket (\mathbf{st}_2, \mathbf{c}_2) = \mathcal{R} (\mathbf{st}_1 \llbracket \mathbf{cs}, (\mathbf{c}_1 \wedge \mathbf{c}_2) \rrbracket \mathbf{st}_2)$$

The synchronisation operator for symbolic traces is simply a symbolic version of the traditional traces operator in [50, p70]. It defines a set of constrained symbolic traces that can arise from synchronisation on the events involving channels in the synchronisation set \mathbf{cs} while performing the traces \mathbf{st}_1 and \mathbf{st}_2 .

Definition 17 (Synchronisation operator - symbolic traces).

$$\mathbf{st}_1 \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \mathbf{st}_2 = \mathbf{st}_2 \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \mathbf{st}_1$$

$$\langle \rangle \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \rangle = \{ (\langle \rangle, \mathbf{c}) \} \quad [\textit{provided } \mathbf{c} \textit{ is satisfiable}]$$

$$\langle \rangle \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \rangle = \emptyset \quad [\textit{otherwise}]$$

$$\langle \rangle \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \mathbf{d}.\beta \rangle = \{ (\langle \mathbf{d}.\beta \rangle, \mathbf{c}) \} \quad [\textit{provided } \mathbf{d} \notin \mathbf{cs}, \mathbf{c} \textit{ is satisfiable}]$$

$$\langle \rangle \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \mathbf{d}.\beta \rangle = \emptyset \quad [\textit{otherwise}]$$

$$\langle \langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1 \rangle \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2 \rangle = \{ \mathbf{cst} : (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \mathbf{st}_2 \bullet (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{cst}.1, \mathbf{cst}.2) \} \\ [\textit{provided } \mathbf{d}_1 \in \mathbf{cs}, \mathbf{d}_2 \notin \mathbf{cs}]$$

$$\langle \langle \mathbf{d}.\alpha \rangle \wedge \mathbf{st}_1 \rangle \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \langle \mathbf{d}.\beta \rangle \wedge \mathbf{st}_2 \rangle = \{ \mathbf{cst} : \mathbf{st}_1 \llbracket \mathbf{cs}, (\mathbf{c} \wedge \alpha = \beta) \rrbracket \mathbf{st}_2 \bullet (\langle \mathbf{d}.\beta \rangle \wedge \mathbf{cst}.1, \mathbf{cst}.2) \} \\ [\textit{provided } \mathbf{d} \in \mathbf{cs}]$$

$$\langle \langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1 \rangle \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2 \rangle = \emptyset \quad [\textit{provided } \mathbf{d}_1 \neq \mathbf{d}_2, \mathbf{d}_1 \in \mathbf{cs}, \mathbf{d}_2 \in \mathbf{cs}]$$

$$\langle \langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1 \rangle \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2 \rangle = \quad [\textit{provided } \mathbf{d}_1 \notin \mathbf{cs}, \mathbf{d}_2 \notin \mathbf{cs}] \\ \{ \mathbf{cst} : (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \mathbf{st}_2 \bullet (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{cst}.1, \mathbf{cst}.2) \} \cup \\ \{ \mathbf{cst} : \mathbf{st}_1 \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \langle \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2 \rangle \bullet (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{cst}.1, \mathbf{cst}.2) \}$$

We, first of all, establish that the synchronisation operator is commutative, and then provide an inductive definition. If both traces are empty, then we get only the empty trace itself, with constraint \mathbf{c} . A proviso guarantees that \mathbf{c} is satisfiable, otherwise no synchronisation traces arise. If there is a conflict in the synchronisation requirements of the traces, then no synchronisation arises either. A conflict occurs when one of the traces requires synchronisation on an event on which the other is not ready to agree (either because it is empty, or its next event is a communication on a different channel also in the synchronisation set.) When there is a synchronisation on events $\mathbf{d}.\alpha$ and $\mathbf{d}.\beta$, we make a(n arbitrary) choice to use the symbolic variable β in the resulting synchronisation trace. This means that α is not used in that trace, but is mentioned in the constraint $\mathbf{c} \wedge \alpha = \beta$. To remove any such symbolic variables that become redundant, in Definition 16, we apply \mathcal{R} to each of the traces in $(\mathbf{st}_1, \mathbf{c}_1) \llbracket \mathbf{cs} \rrbracket (\mathbf{st}_2, \mathbf{c}_2)$.

As already said, the traces in $\mathbf{st}_1 \llbracket \mathbf{cs}, \mathbf{c} \rrbracket \mathbf{st}_2$ are (potentially) over different alphabets. In the definition of $S\text{ctraces}_{\mathbf{cs}}^{\mathbf{a}, (\mathbf{a}_1, \mathbf{a}_2)}(\mathbf{P}_1, \mathbf{P}_2)$ above, the normalisation function \mathcal{N} ensures that the constrained symbolic traces resulting from the synchronisation are over the alphabet \mathbf{a} . Alphabets keep the definition of trace equivalence simple. When defining compositions of (local) traces, however, we have the added complication of requiring normalisation. The normalisation function \mathcal{N} , which is defined below, however, is just a renaming.

Definition 18 (Normalisation). For every $(\mathbf{st}, \mathbf{c})$ and alphabet \mathbf{a} , we define $\mathcal{N}^{\mathbf{a}}(\mathbf{st}, \mathbf{c})$ as follows.

$$\mathcal{N}^{\mathbf{a}}(\mathbf{st}, \mathbf{c}) = (\mathbf{st}[\mathbf{a}/\alpha\mathbf{st}], \mathbf{c}[\mathbf{a}/\alpha\mathbf{st}])$$

This just changes the alphabet of $(\mathbf{st}, \mathbf{c})$, whatever that is, to \mathbf{a} .

The choice of alphabets \mathbf{a}_1 and \mathbf{a}_2 to express the traces of the parallel processes P_1 and P_2 is irrelevant when calculating $Scstraces_{cs}^{\mathbf{a},(\mathbf{a}_1,\mathbf{a}_2)}(P_1, P_2)$ (see Lemma 1 in Appendix C). The alphabet of the traces in this set is always \mathbf{a} . From now on, therefore, we omit these alphabets.

Our main result in this section is given by the following theorem, which establishes that every synchronising trace in $Scstraces_{cs}^{\mathbf{a},(\mathbf{a}_1,\mathbf{a}_2)}(P_1, P_2)$ is a constrained symbolic trace of the parallelism $P_1 \parallel_{cs} P_2$. In addition, they identify all traces of the parallelism that include at least one synchronisation.

Theorem 2.

$$Scstraces_{cs}^{\mathbf{a},(\mathbf{a}_1,\mathbf{a}_2)}(P_1, P_2) = \{ \mathbf{cst} : cstraces^{\mathbf{a}}(P_1 \parallel_{cs} P_2) \mid \mathbf{cst} \upharpoonright_C cs \neq \langle \rangle \}$$

This theorem establishes that if we build tests from $Scstraces_{cs}^{\mathbf{a},(\mathbf{a}_1,\mathbf{a}_2)}(P_1, P_2)$, then we get a subset of the exhaustive test set. Its proof, and those of other theorems introduced here, can be found in Appendix D.

4.3. Compositionality of parallelism

So far, we have, for the sake of clarity and motivation, considered only a parallelism of two processes P_1 and P_2 . It has been observed [35] that in practice the manifestation of a large majority of concurrency bugs involves no more than two threads (which correspond to parallel processes in *Circus*). However, we want to consider arbitrary process expressions; for instance, if we have a parallelism $(P_1 \parallel_{csA} P_2) \parallel_{csB} P_3$, then we may be interested in the synchronisations required by both csA and csB . In this case, one of the parallel components is not a simple process, but a parallelism $(P_1 \parallel_{csA} P_2)$. A compositional approach, in which we build tests based on the (local) traces of P_1 , P_2 , and P_3 , would be ideal.

To discuss this approach, we consider generalisations of the definitions of $ScstraceP_{cs}^{\mathbf{a},(\mathbf{a}_1,\mathbf{a}_2)}(P_1, P_2)$ and $Scstraces_{cs}^{\mathbf{a},(\mathbf{a}_1,\mathbf{a}_2)}(P_1, P_2)$ where the sets of trace pairs or traces are built, not from processes, but from sets of traces $csts_1$ and $csts_2$. The definitions of $ScstraceP_{cs}(\mathbf{csts}_1, \mathbf{csts}_2)$ and $Scstraces_{cs}^{\mathbf{a}}(\mathbf{csts}_1, \mathbf{csts}_2)$ are very similar to those of $ScstraceP_{cs}^{\mathbf{a},(\mathbf{a}_1,\mathbf{a}_2)}(P_1, P_2)$ and $Scstraces_{cs}^{\mathbf{a},(\mathbf{a}_1,\mathbf{a}_2)}(P_1, P_2)$, and are in Appendix A.

The problem that we face is that selection of synchronisation traces is actually not compositional. Precisely, if we consider $(P_1 \parallel_{csA} P_2) \parallel_{csB} P_3$ as suggested above, for instance,

$$\begin{aligned} & Scstraces_{csB}^{\mathbf{a}}(Scstraces_{csA}^{\mathbf{a}_2}(cstraces^{\mathbf{a}_1}(P_1), cstraces^{\mathbf{a}_2}(P_2)), cstraces^{\mathbf{a}_3}(P_3)) \\ & \neq \\ & Scstraces_{csB}^{\mathbf{a}}(cstraces^{\mathbf{a}_2}(P_1 \parallel_{csA} P_2), cstraces^{\mathbf{a}_3}(P_3)) \end{aligned}$$

An example illustrates the issue here.

Example 23. *The (stateless) processes that we consider are as follows.*

$$\begin{aligned} P_1 &= a \longrightarrow \mathbf{Stop} \square b \longrightarrow \mathbf{Stop} \\ P_2 &= a \longrightarrow \mathbf{Stop} \\ P_3 &= b \longrightarrow \mathbf{Stop} \end{aligned}$$

We then consider the specification $P \hat{=} (P_1 \parallel_{\{a\}} P_2) \parallel_{\{b\}} P_3$. First of all, we consider the component $P_1 \parallel_{\{a\}} P_2$, which is itself a parallelism. The traces of this parallelism are $\langle \rangle$, $\langle a \rangle$, and $\langle b \rangle$. Its only synchronisation trace, on the other hand, is $\langle a \rangle$. If we use this synchronisation trace as a basis to construct the synchronisation traces of P , we get the empty set of traces. It, however, has synchronisation trace $\langle b \rangle$, and if we want to cover both synchronisations of the specification, we actually need $\langle a \rangle$ and $\langle b \rangle$. \square

This means that, in general, to construct a test set to cover all synchronisations in a model, we have to construct a test set for each of the parallelisms, considering the whole set of constrained symbolic traces of each of the components, and not only their synchronisation traces. If the synchronisation sets of the parallelisms are the same, however, we do have compositionality, as established by the following theorem.

Theorem 3.

$$\begin{aligned} & Scstraces_{cs}^a(Scstraces_{cs}^{a_{12}}(cstraces^{a_1}(P_1), cstraces^{a_2}(P_2)), cstraces^{a_3}(P_3)) \\ &= \\ & Scstraces_{cs}^a(cstraces^{a_{12}}(P_1 \parallel_{cs} P_2), cstraces^{a_3}(P_3)) \end{aligned}$$

To consider arbitrary process expressions, however, this is not enough.

In general, we need to understand how to construct tests to cover a synchronisation between parallel components occurring in any part of a system model. This is addressed in the next section.

4.4. Compositionality of other operators

We define below the function $selectS^a(P)$, which includes the set of all constrained symbolic traces of P over the alphabet \mathbf{a} that cover at least one synchronisation in P , if any. The set of tests defined from all traces in $selectS^a(P)$ covers all reachable synchronisations of P in all possible ways.

Definition 19 (Selection).

$$\begin{aligned} selectS^a(\text{begin state } [x : T] \bullet \mathbf{A} \text{ end}) &= \emptyset \\ selectS^a(P_1 \parallel_{cs} P_2) &= Scstraces_{cs}^a(P_1, P_2) \\ selectS^a(P_1 \square P_2) &= selectS^a(P_1) \cup selectS^a(P_2) \\ selectS^a(P_1 \sqcap P_2) &= selectS^a(P_1) \cup selectS^a(P_2) \\ selectS^a(P_1 ; P_2) &= selectS^a(P_1) \cup \mathcal{N}^a(\{tcstraces^{a_1}(P_1) \text{ ccat } selectS^{a_2}(P_2)\}) \quad [provided \text{ disjoint}(\mathbf{a}_1, \mathbf{a}_2)] \\ selectS^a(P \setminus cs) &= \mathcal{N}^a(\{selectS^a(P)\} \upharpoonright_{cs}) \end{aligned}$$

A basic process, defined by specifying its state and action, rather than in terms of other processes, has no synchronisation traces. The synchronisation traces of a parallelism $P_1 \parallel_{cs} P_2$ are given by $Scstraces_{cs}^a(P_1, P_2)$. As established by Theorem 3, in selecting traces for a parallelism, we can take advantage of nested parallelisms over the same synchronisation set to optimise the procedure and proceed in a compositional way. This is not considered in the specification of $selectS^a(P)$, but can be exploited in an algorithm.

The synchronisation traces of a sequence $P_1 ; P_2$ include those of P_1 and the continuations of the terminating traces of P_1 that cover the synchronisations of P_2 . To define these, we use the set $tcstraces^a(P)$ of terminating traces of a process P . Its definition below is similar to that of $cstraces^a(P)$; a terminating trace is identified by its leading to a configuration $(c_2 \mid s_2 \models \text{Skip})$ for **Skip**.

Definition 20 (Terminating traces).

$$\begin{aligned} tcstraces^a(\text{begin state } [x : T] \bullet \mathbf{A} \text{ end}) &= tcstraces^a(w_0 \in T, x := w_0, \mathbf{A}) \\ tcstraces^a(c_1, s_1, \mathbf{A}_1) &= \{st, c_2, s_2, \mathbf{A}_2 \mid \alpha st \leq \mathbf{a} \wedge (c_1 \mid s_1 \models \mathbf{A}_1) \xrightarrow{st} (c_2 \mid s_2 \models \text{Skip}) \bullet \mathcal{R}(st, c_2)\} \end{aligned}$$

To cover the synchronisations in P_2 , we need to prefix its synchronisation traces with terminating traces of P_1 . We use the operator $csts_1 \text{ ccat } csts_2$; the traces in $csts_1 \text{ ccat } csts_2$ are those formed by concatenating all pairs of symbolic traces and conjoining their associated constraints. Since, in general, the length of the traces in $tcstraces^{a_1}(P_1)$ vary, it is not possible to predict the alphabet of the traces in $tcstraces^{a_1}(P_1) \text{ ccat } selectS^{a_2}(P_2)$. In general, different traces in this set may have different alphabets. We, therefore, use in Definition 19 the normalisation function \mathcal{N} to obtain a set of traces over the alphabet \mathbf{a} .

For a hiding $P \setminus cs$, since we are interested in covering the hidden synchronisations, we consider the synchronisation traces of P . We remove from them, however, the events that represent communications through the channels in cs . For that, we use the projection function $cst \upharpoonright_{cs}$, whose resulting constrained symbolic trace contains only the events of cst that are not communications over channels in cs . The alphabet of a symbolic trace resulting from the application of this operator is not predictable, so we again use the normalisation function to obtain in $selectS^a(P \setminus cs)$ a set of constrained symbolic traces over \mathbf{a} .

Example 24. Both synchronising traces in Example 20 give rise to the synchronising trace $((\text{in.}\alpha), \text{true})$ of Protocol. \square

Our main result here is introduced below: our selection strategy identifies a subset of the traces of a process.

Theorem 4.

$$\text{selectS}^a(\mathbb{P}) \subseteq \text{cstraces}^a(\mathbb{P})$$

This means that selection based on $\text{selectS}^a(\mathbb{P})$ identifies an unbiased test set (that is, a set of traces that can be used to construct an unbiased test set). In terms of the framework defined in the previous section, the property $\pi(t, P)$ that can be used to achieve selection based on synchronisation coverage is $t \in \text{selectS}^a(\mathbb{P})$. We consider the use of $\text{selectS}^a(\mathbb{P})$ for selection of tests for traces refinement next.

4.5. More selective synchronisation coverage

With the property $\pi(t, SP) = t \in \text{selectS}^a(SP)$, we can define, using our framework from Section 3, two exhaustive test sets: one for existential and one for universal selection based on traces. In this example, the traces and tests are symbolic. We use the framework in Section 3, but with definitions adapted (in the expected and straightforward way) to consider symbolic traces and tests. The symbolic exhaustive test sets are shown below. In these definitions, we also rely on Theorem 4 to conclude that $t \in \text{selectS}^a(SP)$ implies that $t \in \text{cstraces}(SP)$ and simplify the definitions obtained directly using (1) and (2).

$$\begin{aligned} \text{SyncExist}_T^{\text{trace}}(SP) &= \{ \text{AllTests}_T(\text{cst}, SP) \mid \text{cst} \in \text{selectS}^a(SP) \} \\ \text{SyncUniv}_T^{\text{trace}}(SP) &= \{ \text{ST}_T(\text{cst}, \text{cse}) \mid \text{cst} \in \text{selectS}^a(SP) \wedge \text{cse} \in \overline{\text{csinitials}}(SP, \text{cst}) \} \end{aligned}$$

Existential selection, in this example, is not very interesting, since its goal is to run only one test covering at least one synchronisation. Universal selection, on the other hand, is the basis for the property: “all tests that involve at least one synchronisation in the trace are attempted”.

The problem that we face in this section is the fact that the above sets of symbolic tests are still infinite, and, in addition, each symbolic test has an infinite number of instantiations as well. We consider next how to obtain a finite set of symbolic tests for universal selection.

For selection of symbolic tests, what we suggest is the use of a regularity hypothesis. If the model does not have unbounded nondeterminism, a regularity hypothesis that limits the size of the symbolic traces defines a finite set of symbolic tests. It is well known, however, that the choice of the limiting size must be careful, to avoid throwing away too many interesting tests [18]. In the case of synchronisation coverage, it is sensible to limit the size on the traces in such a way that ensures that all channels are used at least once.

Example 25. For synchronisation tests for the parallelism between the processes *Components* and *Medium*, the bound should be at least 8, to ensure that traces with events representing communications on *out* are included. For this system, we define the symbolic exhaustive set of tests (with respect to synchronisation coverage and our regularity hypothesis), as follows.

$$\begin{aligned} \text{FSyncUniv}_T^{\text{trace}}(\text{Components} \llbracket \{ \text{left}, \text{right} \} \rrbracket \text{Medium}) = \\ \{ \text{ST} \mid \text{ST} \in \text{SyncUniv}_T^{\text{trace}}(\text{Components} \llbracket \{ \text{left}, \text{right} \} \rrbracket \text{Medium}) \wedge \# \text{ST} \leq 9 \} \end{aligned}$$

We use the operator $\# \text{ST}$ to denote the size of a symbolic test: the size of its symbolic trace plus one. \square

For selection of instances of symbolic tests, what we suggest is the use of uniformity hypotheses based on the constraint of the symbolic traces and of the symbolic forbidden continuations. We can use the constraints to take into account information about the updates of data specified in a symbolic trace; an extreme decomposition uses a DNF characterisation of the constraints to define several domains. To be more selective, we consider each constraint as defining a single domain. In this case, selection is based on the interactions and paths of execution in the model defined by the operational semantics used to generate the symbolic traces. This means that we take the existential selection of instances of the symbolic tests (discussed in Section 3.4). This runs, if possible, one instance of every selected symbolic test.

Example 26. *The driver for existential selection of instances of the tests in Example 25 is as follows.*

$$\begin{aligned} & \text{SyncUnivExistDriver}^{\text{trace}}(\text{Components } \llbracket \{ \text{left}, \text{right} \} \rrbracket \text{Medium}, \text{SUT}) = \\ & \left(\left\| \left\| \left\| ST : \text{FSyncUniv}_T^{\text{trace}}(\text{Components } \llbracket \{ \text{left}, \text{right} \} \rrbracket \text{Medium}) \bullet \right. \right. \right. \\ & \left. \left. \left. \left(\left\| \left\| T : \text{instTest}(ST) \bullet \text{Execution}_{\text{SUT}}^{\text{SP}}(T) \right\| \llbracket \{ \text{inc}, \text{pass}, \text{fail} \} \rrbracket \text{TMonitor} \right. \right. \right. \right. \end{aligned}$$

For each of the finitely many symbolic tests in $\text{FSyncUniv}_T^{\text{trace}}(\text{Components } \llbracket \{ \text{left}, \text{right} \} \rrbracket \text{Medium})$, we have the independent execution of its instances under the control of a *TMonitor*. As soon as an instance reaches a conclusive verdict, other instances are discarded. \square

We observe that, as illustrated in the above example, the search for a single instance may still not terminate, if none is implemented. Static choice of an instance is not a good idea, as extensively discussed in Section 3.

The exhaustivity of our test sets rely on the complete test coverage assumption to guarantee, when the SUT is nondeterministic, the execution of all its behaviours. In the case of parallel composition, it implies the spontaneous coverage of all synchronisation-race variants: this assumption can be weakened by using specific schedulers when driving the tests, such as in reachability testing [34], or CHESS [1].

5. Related works

One of the first well known coverage criteria for concurrent programs was developed by Taylor et al. [55] for Ada. It is based on notions of concurrency states and concurrency graphs used to define classical-graph coverage criteria (all-concurrency-paths, all-edges-between-cc-states, and all-cc-states), and a weaker criterion called *all-possible-rendezvous*, which requires that for all concurrency states that involve a rendezvous, there is at least one path through this state. This notion of coverage hides the sequential activities in the concurrency states, thus focusing on synchronisation anomalies. The authors also sketch the idea of forcing particular concurrency paths when driving test execution. As mentioned in Section 4, synchronisation coverage for Ada programs has been more precisely defined and implemented by Bron et al. [5].

Carver, Lei, Tai and others has developed reachability testing [7, 34], which provides ways to execute deterministically a concurrent program and reproduce a given test case. It is a purely dynamic method for generating *synchronisation sequences* (SYN-sequences) on the fly. Reachability testing algorithms are exhaustive; this ensures the execution of all the possible SYN-sequences for a given input, provided there is no unbounded loop. In this case, a bound on the length of the SYN-sequences is introduced.

For synchronous message passing, like in *Circus*, SYN-sequences contain pairs of sending and receiving events exercised in an execution. Deterministic execution is led according to a given SYN-sequence. Reachability testing does not address selection of test inputs or SYN-sequences, but specification-based reachability testing is addressed in [8] with a limited specification language for expressing local constraints on the SYN-sequences. When the only source of nondeterminism is race variants, reachability testing can be used to ensure the complete test assumption used in our work, since it ensures execution of all SYN-sequences.

The use of a reachability graph for structural test-sequence generation for concurrent programs is reported in [61]. A more recent approach by Souza et al. [54] reports some experiments with reachability testing and a few coverage criteria based on a notion of *sync-edge association*, which seems similar to the coverage of those edges corresponding to a synchronisation pair. This is close to the notion of synchronisation coverage that we proposed for *Circus*, but it is based on static analysis of the program.

The next generation of reachability testing algorithms is embedded in CHESS [39, 1]. This tool repeatedly runs concurrent tests ensuring that every run takes a different interleaving. If an interleaving results in an error, CHESS can reproduce it for improved debugging. The only perturbation introduced is a thin wrapper layer between the SUT and the concurrency API. The wrappers are similar to implementations of our drivers, but just like in all reachability testing techniques, they are not concerned with selection.

Dynamic test generation [6] avoids infeasible paths when generating tests from the text of a program and covering all paths. This is achieved by symbolic and standard execution in an integrated way. Test generation starts by running the program on an arbitrary concrete test input, and recording the choices and

the path followed. Afterwards, other tests are generated by switching choices to execute another path: a constraint solver is used to check that the changes lead to feasible paths. This method has been mainly applied to testing based on the source code of programs, as in PathCrawler [60], CUTE [53], and Dart [25].

A first difference between these approaches and our work is that our models are potentially more abstract than the source code of programs. Secondly, similarity comes from the need to attempt alternative paths when a test is not possible, but the main difference is that the SUT is not required to implement all the traces of our models. Thus, our concern is absence of specification traces rather than unsatisfiability, so we do not use constraint solving. We note that in the presence of loops, dynamic test generation may not terminate. Termination is ensured by introducing some limit to the length of paths, which corresponds to an implicit regularity hypothesis on the SUT. Similarly, our drivers for symbolic tests may not terminate.

Adaptive testing has been studied for the derivation of elaborate test experiments from nondeterministic finite state machines [33] for possibly nondeterministic implementations [46, 27, 29]. The idea is to take into account the output of the SUT to determine how to proceed with the experiment. It means that there is no preset input sequence, but a rooted tree that anticipates all possible sequences. It is very similar to the factorised tree-shaped tests we present in Section 3.3, which lets the SUT to lead the test. The role of our various test drivers is a generalisation of this approach, where, in addition to the possibility of non-implemented traces, we introduce a variety of selection strategies.

For a deterministic SUT, further work [30] has considered also the optimisation of a testing experiment by ordering the adaptive tests. In our setting, we can consider, for example, that, if the SUT is deterministic, an inconclusive test verdict indicates that all tests for the same trace or any of its extensions is also going to produce an inconclusive verdict. Therefore, they are useless and can be eliminated. Sophisticated drivers are required to order the tests and perform such optimisations.

An important feature that is missing in the semantics of *Circus* (and CSP) is a distinction of inputs and outputs. The difference in the controllability of these events, with inputs controlled by the environment and outputs by the process, is clearly important for testing. The work in [16] discusses this point in the context of CSP and, in particular, the impact of a refinement relation that captures the distinction between inputs and outputs on its testing theory. We note, however, that the distinction between inputs and outputs affects the refusals of a process, but not its traces. For this reason, the results presented here for traces refinement are not affected by considerations of controllability.

Peleska and Siegel [44, 45] applied CSP-based testing and developed a tool and technique for automation. More recently, CSP was used by Nogueira et al. to formalise a notion of conformance traditionally associated with input-output labelled transition systems [41], with refinement model checking used for generation and selection of tests. This work provides a technique and a tool for generation and selection of tests based on the use of FDR, the CSP refinement model checker. In both these works, universal selection strategies are used, and, therefore, the issue on non-implemented traces does not arise.

In [37] Massink et al. have studied testing based on statecharts. A testing theory, in the sense of [40], that is, for models rather than systems, as well as a notion of conformance and a test-case generation algorithm, were presented. The theoretical framework developed follows a pattern very similar to ours, but the considered language and the conformance relation are very different. Exhaustiveness of the generated test set is proved, but test selection strategies are left as future work.

A symbolic version of the *ioco* relation [57] was stated by Frantzen et al. in [22] for Input-Output Symbolic Transition Systems (IOSTS) implemented by input-enabled IOSTS. There is a slightly different notion of constrained symbolic extended traces, with two constraints, one on the interaction variables, and another on the updates of the state variables. Location coverage is mentioned as a perspective: a sort of symbolic state coverage. This would raise the same issues we have studied and solved here, since it is an existential coverage criterion, and *ioco* does not require all the traces to be implemented.

6. Conclusions

In this paper, we have identified and illustrated the problem of selection of tests for establishing traces refinement. In short, it is not possible to have a static selection, because a correct SUT may not implement the selected traces, in which case coverage ensured statically may not be achieved during experiments.

As a solution, we have proposed the use of strategies that combine the definition of an exhaustive test set with respect to the selection criteria of interest and a driver. Although, we have considered the main kinds of selection criteria, it is not ever possible to claim that we have catered for every imaginable criterion. We believe, however, that the approaches that we have proposed and illustrated for comprehensive categories of criteria can be usefully combined or extended to cater for sophisticated selection criteria. We have, in this paper, considered selection of both concrete and symbolic tests.

We have considered criteria based on one or several (multiple) properties of traces or of tests, and that require one (existential) or all (universal) implemented traces or tests that satisfy the property(ies) to be run. For each such kind of criteria, we have formalised the construction of the exhaustive test set and described a driver using *Circus* constructs. These driver specifications provide an accurate description of the requirements of the test experiments, but practical techniques have to consider efficient use of the SUT and of constraint solvers to avoid the use of too many copies of the SUT or instantiations of symbolic tests.

To illustrate the use of our framework, we have formalised a selection criterion for testing based on parallel *Circus* specifications. It is appropriate for integration testing, and requires coverage of synchronisations between parallel processes in a specification. We have formally defined this trace selection criterion and, specified an algorithm for selection. We have proved that this algorithm generates traces of the specification, and, therefore, unbiased tests. This is used to define a finite set of symbolic tests and accompanying driver.

The problem raised by the possibility of selecting non-implemented traces is also raised by a different conformance relation that we have studied previously in the context of *Circus*, namely, *conf*. What *conf* requires is that, for traces that are common to the specification and the SUT, the SUT deadlocks only when a deadlock is allowed in the model. As opposed to the tests for traces refinement, tests for *conf* are positive. They check that, after a given trace of the specification, at least one event from a set of minimal acceptances is accepted by the SUT. Since *conf* is restricted to traces that are common to the SUT and the specification, if the trace is not implemented by the SUT, we again have an inconclusive verdict. We show in [14] that the same solutions proposed here are applicable in testing for *conf*.

This is an important result because we have previously shown that failures refinement, another important conformance relation for both CSP and *Circus*, can be characterised as the conjunction of traces refinement and *conf* [10, 13]. An exhaustive test set for failures refinement is, therefore, the union of the exhaustive test sets for traces refinement and for *conf*. Moreover, while traces refinement is concerned with safety properties and cannot, for instance, detect a mismatch in the alphabet of events of an SUT, *conf* can be used to identify such a problem. If the SUT insists in performing an event outside of the alphabet of the specification, unless a deadlock is allowed by the specification, a test for *conf* can identify the problem.

A popular conformance relation is *ioco* [57]. Since it allows that some traces of the specification not to be implemented, our results are of interest for test selection in the case of *ioco*.

As for future work we are interested in performing case studies, both in the form of the definition of additional selection criteria and experiments of testing actual systems. In that respect, the automation of the *Circus* testing theory presented in [19, 20] is of great relevance.

Also of interest is the application of the ideas presented here to the notion of traces for *Circus* defined in [12] and [15]. In that work, we consider traces that record information about the data operations (specification statements, assignments, guards, and so on) in the *Circus* specification, and support selection techniques based on the structure of the model: in [15] we present several data flow coverage criteria for *Circus* specifications. Just like the symbolic traces of the original *Circus* testing theory, these symbolic traces do not need to be implemented by an SUT. So, the same issues raised here apply.

Acknowledgements. We are grateful to the Royal Society and the CNRS for funding our collaboration via their International Exchange scheme, and to Rob Hierons for various discussions. We are also indebted to anonymous referees for their careful revision and useful suggestions.

References

- [1] T. Ball, S. Burckhardt, K. E. Coons, M. Musuvathi, and S. Qadeer. Preemption sealing for efficient concurrency testing. In *16th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *Lecture Notes in Computer Science*, pages 420 – 434. Springer, 2010.

- [2] G. Barrett. Model checking in practice: The T9000 Virtual Channel Processor. *IEEE Transactions on Software Engineering*, 21(2):69 – 78, 1995.
- [3] L. Bentakouk, P. Poizat, and F. Zaidi. Checking the behavioral conformance of web services with symbolic testing and an smt solver. In Springer, editor, *Tests and Proofs*, volume 6706 of *Lecture Notes in Computer Science*, pages 33 – 50, 2011.
- [4] G. Bernot, M.-C. Gaudel, and B. Marre. Software testing based on formal specifications: a theory and a tool. *Software Engineering Journal*, 6(6):387 – 405, 1991.
- [5] A. Bron, E. Farchi, Y. Magid, Y. Nir, and S. Ur. Applications of synchronization coverage. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 206 – 212, 2005.
- [6] C. Cadar, P. Godefroid, S. Khurshid, C. S. Păsăreanu, K. Sen, N. Tillmann, and W. Visser. Symbolic execution for software testing in practice: preliminary assessment. In *33rd International Conference on Software Engineering*, pages 1066 – 1071. ACM, 2011.
- [7] R. H. Carver and K.-C. Tai. Replay and Testing for Concurrent Programs. *IEEE Software*, 8(2):66 – 74, 1991.
- [8] R. H. Carver and Kuo-Chung Tai. Use of Sequencing Constraints for Specification-Based Testing of Concurrent Programs. *IEEE Transactions on Software Engineering*, 24:471 – 490, 1998.
- [9] A. L. C. Cavalcanti, P. Clayton, and C. O'Halloran. From Control Law Diagrams to Ada via *Circus*. *Formal Aspects of Computing*, 23(4):465 – 512, 2011.
- [10] A. L. C. Cavalcanti and M.-C. Gaudel. Testing for Refinement in CSP. In *9th International Conference on Formal Engineering Methods*, volume 4789 of *Lecture Notes in Computer Science*, pages 151 – 170. Springer-Verlag, 2007.
- [11] A. L. C. Cavalcanti and M.-C. Gaudel. A note on traces refinement and the *conf* relation in the Unifying Theories of Programming. In A. Butterfield, editor, *Unifying Theories of Programming 2008*, volume 5713 of *Lecture Notes in Computer Science*. Springer-Verlag, 2010.
- [12] A. L. C. Cavalcanti and M.-C. Gaudel. Specification Coverage for Testing in *Circus*. In S. Qin, editor, *Unifying Theories of Programming*, volume 6445 of *Lecture Notes in Computer Science*, pages 1 – 45. Springer-Verlag, 2010.
- [13] A. L. C. Cavalcanti and M.-C. Gaudel. Testing for Refinement in *Circus*. *Acta Informatica*, 48(2):97 – 147, 2011.
- [14] A. L. C. Cavalcanti and M.-C. Gaudel. A note on test selection for *conf* refinement. Rapport LRI 1569, Université de Paris-Sud, 2013. Available at www.lri.fr/bibli/Rapports-internes/2013/RR1569.pdf.
- [15] A. L. C. Cavalcanti and M.-C. Gaudel. Data Flow coverage for *Circus*-based testing. In *Fundamental Approaches to Software Engineering*, Lecture Notes in Computer Science, 2014.
- [16] A. L. C. Cavalcanti and R. M. Hierons. Testing with Inputs and Outputs in CSP. In *Fundamental Approaches to Software Engineering*, volume 7793 of *Lecture Notes in Computer Science*, pages 359 – 374, 2013.
- [17] A. L. C. Cavalcanti, A. C. A. Sampaio, and J. C. P. Woodcock. A Refinement Strategy for *Circus*. *Formal Aspects of Computing*, 15(2 - 3):146 – 181, 2003.
- [18] R. K. Dong and P. G. Frankl. The ASTOOT approach to testing object-oriented programs. *ACM Transactions on Software Engineering and Methodology*, 3(2):103 – 130, 1994.
- [19] A. Feliachi. *Semantics-Based Testing for Circus*. PhD thesis, LRI, Université de Paris-Sud, 2012.
- [20] A. Feliachi, M. C. Gaudel, M. Wenzel, and B. Wolff. The *Circus* Testing Theory Revisited in Isabelle/HOL. In L. Groves and J. Sung, editors, *15th International Conference on Formal Engineering Methods*, volume 8144 of *Lecture Notes in Computer Science*, pages 243 – 260. Springer, 2013.
- [21] C. Fischer. How to Combine Z with a Process Algebra. In J. Bowen, A. Fett, and M. Hinchey, editors, *The Z Formal Specification Notation*. Springer-Verlag, 1998.
- [22] L. Frantzen, J. Tretmans, and T. A. C. Willemse. A Symbolic Framework for Model-Based Testing. In K. Havelund, M. Núñez, G. Rosu, and B. Wolff, editors, *Formal Approaches to Software Testing and Runtime Verification*, number 4262 in *Lecture Notes in Computer Science*, pages 40 – 54. Springer, 2006.
- [23] L. Freitas and J. P. McDermott. Formal methods for security in the xenon hypervisor. *International Journal on Software Tools for Technology Transfer*, 13(5):463 – 489, 2011.
- [24] S. Fujiwara and G. Bochmann. Testing non-deterministic finite state machines with fault coverage. In *4th International Workshop on Protocol Test Systems*, pages 267–280. North-Holland Publishing Co., 1991.
- [25] P. Godefroid, N. Klarlund, and K. Sen. DART: directed automated random testing. In *ACM SIGPLAN 2005 Conference on Programming Language Design and Implementation, PLDI*, pages 213 – 223. ACM, 2005.
- [26] A. Hall and R. Chapman. Correctness by construction: Developing a commercial secure system. *IEEE Software*, 19(1):18 – 25, 2002.
- [27] R. M. Hierons. Adaptive testing of a deterministic implementation against a nondeterministic finite state machine. *Computer Journal*, 41(5):349 – 355, 1998.
- [28] R. M. Hierons. Comparing test sets and criteria in the presence of test hypotheses and fault domains. *ACM Transactions on Software Engineering and Methodology*, 11(4):427–448, 2002.
- [29] R. M. Hierons. Applying adaptive test cases to nondeterministic implementations. *Information Processing Letters*, 98(2):56 – 60, 2006.
- [30] R. M. Hierons and H. Ural. Reducing the cost of applying adaptive test cases. *Computer Networks*, 51(1):224 – 238, 2007.
- [31] E. Itoh, Y. Kawaguchi, Z. Furukawa, and K. Ushijima. Ordered Sequence Testing Criteria for Concurrent Programs and the Support Tool. In *1st Asia-Pacific Software Engineering Conference*, pages 236 – 245. IEEE, 1994.
- [32] K. G. Larsen, M. Mikucionis, B. Nielsen, and A. Skou. Testing real-time embedded software using uppaal-tron: an industrial case study. In *5th ACM International Conference on Embedded Software*, pages 299 – 306. ACM, 2005.
- [33] D. Lee and M. Yannakakis. Principles and methods of testing finite state machines - A survey. In *Proceedings of the IEEE*, volume 84, pages 1090 – 1126, 1996.

- [34] Y. Lei and R. H. Carver. Reachability testing of concurrent programs. *IEEE Transactions on Software Engineering*, 32(6):382 – 403, 2006.
- [35] S. Lu, S. Park, E. Seo, and Y. Zhou. Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. In *International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 329 – 339, 2008.
- [36] B. P. Mahony and J. S. Dong. Blending Object-Z and Timed CSP: An Introduction to TCOZ. In *The 20th International Conference on Software Engineering (ICSE'98)*, pages 95 – 104. IEEE Computer Society Press, 1998.
- [37] M. Massink, D. Latella, and S. Gnesi. On testing UML statecharts. *The Journal of Logic and Algebraic Programming*, 69(1-2):1 – 74, 2006.
- [38] A. Miyazawa and A. L. C. Cavalcanti. Refinement-oriented models of Stateflow charts. *Science of Computer Programming*, 77(10 – 11):1151 – 1177, 2012.
- [39] M. Musuvathi, S. Qadeer, T. Ball, G. Basler, P. Arumuga Nainar, and I. Neamtiu. Finding and reproducing heisenbugs in concurrent programs. In R. Draves and R. van Renesse, editors, *8th USENIX Symposium on Operating Systems Design and Implementation*, pages 267 – 280. USENIX Association, 2008.
- [40] R. De Nicola and M. C. B. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 3(1-2):83 – 133, 1984.
- [41] S. Nogueira, A. C. A. Sampaio, and A. C. Mota. Guided Test Generation from CSP Models. In J. S. Fitzgerald, A. E. Haxthausen, and H. Yenigün, editors, *5th International Colloquium on Theoretical Aspects of Computing*, volume 5160 of *Lecture Notes in Computer Science*, pages 258 – 273. Springer, 2008.
- [42] M. V. M. Oliveira. *Formal Derivation of State-Rich Reactive Programs Using Circus*. PhD thesis, University of York, 2006.
- [43] M. V. M. Oliveira, A. L. C. Cavalcanti, and J. C. P. Woodcock. A UTP Semantics for *Circus*. *Formal Aspects of Computing*, 21(1-2):3 – 32, 2009.
- [44] J. Peleska and M. Siegel. Test automation of safety-critical reactive systems. In *Formal Methods Europe, Industrial Benefits and Advances in Formal Methods*, volume 1051 of *Lecture Notes in Computer Science*, 1996.
- [45] J. Peleska and M. Siegel. Test automation of safety-critical reactive systems. *South African Computer Journal*, 19:53 – 77, 1997.
- [46] A. Petrenko, N. Yevtushenko, A. Lebedev, and A. Das. Nondeterministic state machines in protocol conformance testing. In *6th International Workshop on Protocol Test systems VI*, pages 363 – 378. North-Holland Publishing Co., 1994.
- [47] C. Robinson-Mallett, R. M. Hierons, and P. Liggesmeyer. Achieving communication coverage in testing. *ACM SIGSOFT Software Engineerig Notes*, 31(6):1 – 10, 2006.
- [48] C. Robinson-Mallett, R. M. Hierons, J. Poore, and P. Liggesmeyer. Using communication coverage criteria and partial model generation to assist software integration testing. *Software Quality Control*, 16(2):185 – 211, 2008.
- [49] M. Roggenbach. CSP-CASL: a new integration of process algebra and algebraic specification. *Theoretical Computer Science*, 354(1):42 – 71, 2006.
- [50] A. W. Roscoe. *The Theory and Practice of Concurrency*. Prentice-Hall Series in Computer Science. Prentice-Hall, 1998.
- [51] A. W. Roscoe. *Understanding Concurrent Systems*. Texts in Computer Science. Springer, 2011.
- [52] S. Schneider and H. Treharne. Communicating B Machines. In D. Bert, J. Bowen, M. Henson, and K. Robinson, editors, *ZB'2002: Formal Specification and Development in Z and B*, volume 2272 of *Lecture Notes in Computer Science*, pages 416 – 435, 2002.
- [53] K. Sen, D. Marinov, and G. Agha. CUTE: a concolic unit testing engine for C. In *10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering*, pages 263 – 272. ACM, 2005.
- [54] S. R. S. Souza, P. S. L. Souza, M. C. C. Machado, M. S. Camillo, A. S. Simão, and E. Zaluska. Using Coverage and Reachability Testing to Improve Concurrent Program Testing Quality. In *23rd International Conference on Software Engineering and Knowledge Engineering*, pages 207 – 212, 2011.
- [55] R. N. Taylor, D. L. Levine, and C. D. Kelly. Structural testing of concurrent programs. *IEEE Transactions on Software Engineering*, pages 206 – 215, 1992.
- [56] N. Tillmann, J. de Halleux, and W. Schulte. Parameterized Unit Testing with Pex: Tutorial. In P. H. M. Borba, A. L.C. Cavalcanti, A. C. A. Sampaio, and J. C. P. Woodcock, editors, *Testing Techniques in Software Engineering, Second Pernambuco Summer School on Software Engineering*, volume 6153 of *Lecture Notes in Computer Science*, pages 141 – 202. Springer, 2010.
- [57] J. Tretmans. Test Generation with Inputs, Outputs, and Quiescence. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 1055 of *Lecture Notes in Computer Science*, pages 127 – 146. Springer-Verlag, 1996.
- [58] J. Tretmans and E. Brinksma. TorX: Automated Model Based Testing. In *1st European Conference on Model-Driven Software Engineering*, pages 13 – 25, 2003.
- [59] M. Veanes, C. Campbell, W. Schulte, and P. Kohli. On-the-fly testing of reactive systems. Technical Report MSR-TR-2005-03, Microsoft Research, 2005.
- [60] N. Williams, B. Marre, P. Mouy, and M. Roger. PathCrawler: Automatic generation of path tests by combining static and dynamic analysis. In *5th European Dependable Computing Conference*, volume 3463 of *Lecture Notes in Computer Science*, pages 281 – 292. Springer, 2005.
- [61] W. E. Wong and Y. Lei. Reachability graph-based test sequence generation for concurrent programs. *International Journal of Software Engineering and Knowledge Engineering*, 18(6):803 – 822, 2008.
- [62] J. C. P. Woodcock and J. Davies. *Using Z—Specification, Refinement, and Proof*. Prentice-Hall, 1996.

Appendix A. Formal definitions of operators

Definition 21. We define a pair of alphabets \mathbf{a}_1 and \mathbf{a}_2 to be disjoint if their sets of elements are disjoint: $\text{disjoint}(\mathbf{a}_1, \mathbf{a}_2) = \text{ran } \mathbf{a}_1 \cap \text{ran } \mathbf{a}_2 = \emptyset$.

Definition 22. For every constrained symbolic trace (st, \mathbf{c}) and synchronisation channel set cs , we define the projection $(\text{st}, \mathbf{c}) \upharpoonright_{\mathbf{C}} \text{cs}$ of the trace (st, \mathbf{c}) to channels in cs as $(\text{st}, \mathbf{c}) \upharpoonright_{\mathbf{C}} \text{cs} = \text{st} \upharpoonright_{\mathbf{C}} \text{cs}$, where the similar projection operator for symbolic traces is defined inductively as follows.

$$\begin{aligned} \langle \rangle \upharpoonright_{\mathbf{C}} \text{cs} &= \langle \rangle \\ \langle \langle \mathbf{d}.\alpha \rangle \wedge \text{st} \rangle \upharpoonright_{\mathbf{C}} \text{cs} &= \langle \mathbf{d} \rangle \wedge (\text{st} \upharpoonright_{\mathbf{C}} \text{cs}), \text{ if } \mathbf{d} \in \text{cs} \\ \langle \langle \mathbf{d}.\alpha \rangle \wedge \text{st} \rangle \upharpoonright_{\mathbf{C}} \text{cs} &= \text{st} \upharpoonright_{\mathbf{C}} \text{cs}, \text{ if } \mathbf{d} \notin \text{cs} \end{aligned}$$

Definition 23. For every symbolic trace st and synchronisation channel set cs , we define inductively the projection $\text{st} \upharpoonright_{\mathbf{E}} \text{cs}$ of st to events representing communications over channels in cs as follows.

$$\begin{aligned} \langle \rangle \upharpoonright_{\mathbf{E}} \text{cs} &= \langle \rangle \\ \langle \langle \mathbf{d}.\alpha \rangle \wedge \text{st} \rangle \upharpoonright_{\mathbf{E}} \text{cs} &= \langle \mathbf{d}.\alpha \rangle \wedge (\text{st} \upharpoonright_{\mathbf{E}} \text{cs}), \text{ if } \mathbf{d} \in \text{cs} \\ \langle \langle \mathbf{d}.\alpha \rangle \wedge \text{st} \rangle \upharpoonright_{\mathbf{E}} \text{cs} &= \text{st} \upharpoonright_{\mathbf{E}} \text{cs}, \text{ if } \mathbf{d} \notin \text{cs} \end{aligned}$$

Definition 24. For sets of constrained symbolic traces csts_1 and csts_2 with disjoint alphabets, and channel set cs , we define the set of synchronising pairs $\text{ScstraceP}_{\text{cs}}(\text{csts}_1, \text{csts}_2)$ as follows.

$$\begin{aligned} \text{ScstraceP}_{\text{cs}}(\text{csts}_1, \text{csts}_2) &= \\ &\{ \text{cst}_1 : \text{csts}_1; \text{cst}_2 : \text{csts}_2 \mid (\text{cst}_1 \upharpoonright_{\mathbf{C}} \text{cs}) \neq \langle \rangle \wedge (\text{cst}_2 \upharpoonright_{\mathbf{C}} \text{cs}) \neq \langle \rangle \wedge \text{satisfiable}_{\text{cs}}(\text{cst}_1, \text{cst}_2) \} \end{aligned}$$

Definition 25. For sets of constrained symbolic traces csts_1 and csts_2 , channel set cs , and alphabet \mathbf{a} we define the set $\text{Scstraces}_{\text{cs}}^{\mathbf{a}}(\text{csts}_1, \text{csts}_2)$ of synchronising constrained symbolic traces over \mathbf{a} as follows.

$$\text{Scstraces}_{\text{cs}}^{\mathbf{a}}(\text{csts}_1, \text{csts}_2) = \bigcup \{ \mathbf{p} : \text{ScstraceP}_{\text{cs}}(\text{csts}_1, \text{csts}_2) \bullet \mathcal{N}^{\mathbf{a}}(\mathbf{p}.1 \llbracket \text{cs} \rrbracket \mathbf{p}.2) \}$$

Definition 26.

$$\text{csts}_1 \text{ ccat } \text{csts}_2 = \{ \text{st}_1, \text{st}_2, \mathbf{c}_1, \mathbf{c}_2 \mid (\text{st}_1, \mathbf{c}_1) \in \text{csts}_1 \wedge (\text{st}_2, \mathbf{c}_2) \in \text{csts}_2 \bullet (\text{st}_1 \wedge \text{st}_2, \mathbf{c}_1 \wedge \mathbf{c}_2) \}$$

Definition 27. For every constrained symbolic trace (st, \mathbf{c}) and synchronisation channel set cs , we define the projection $(\text{st}, \mathbf{c}) \downharpoonright_{\mathbf{C}} \text{cs}$ of the trace (st, \mathbf{c}) to channels not in cs as $(\text{st}, \mathbf{c}) \downharpoonright_{\mathbf{C}} \text{cs} = \mathcal{R}(\text{st} \upharpoonright_{\mathbf{C}} \text{cs}, \mathbf{c})$, where the projection operator for symbolic traces is defined inductively as follows.

$$\begin{aligned} \langle \rangle \downharpoonright_{\mathbf{C}} \text{cs} &= \langle \rangle \\ \langle \langle \mathbf{d}.\alpha \rangle \wedge \text{st} \rangle \downharpoonright_{\mathbf{C}} \text{cs} &= \langle \mathbf{d}.\alpha \rangle \wedge (\text{st} \downharpoonright_{\mathbf{C}} \text{cs}), \text{ if } \mathbf{d} \notin \text{cs} \\ \langle \langle \mathbf{d}.\alpha \rangle \wedge \text{st} \rangle \downharpoonright_{\mathbf{C}} \text{cs} &= \text{st} \downharpoonright_{\mathbf{C}} \text{cs}, \text{ if } \mathbf{d} \in \text{cs} \end{aligned}$$

Appendix B. Properties of the Circus operational semantics

The following propositions establish basic properties of the *Circus* operational semantics that we need. A simple property is captured by the proposition below. It states that, if the action does not involve a variable x , or more generally, any of the variables in a state \mathbf{s}_3 , then its value in the state is not changed.

Proposition 1.

$$(\mathbf{c}_1 \mid \mathbf{s}_1 \models \mathbf{A}_1) \xrightarrow{\text{st}} (\mathbf{c}_2 \mid \mathbf{s}_2 \models \mathbf{A}_2)$$

if, and only if,

$$(\mathbf{c}_1 \mid \mathbf{s}_1; \mathbf{s}_3 \models \mathbf{A}_1) \xrightarrow{\text{st}} (\mathbf{c}_2 \mid \mathbf{s}_2; \mathbf{s}_3 \models \mathbf{A}_2)$$

provided $FV(\mathbf{s}_3) \cap FV(\mathbf{A}_1) = \emptyset$.

$FV(\mathbf{A})$ is the set of free variables of the action \mathbf{A} as defined in [42]. In addition, if a constraint \mathbf{c}_3 is over

symbolic variables that are not used in the state nor in the symbolic trace, then it can be ignored.

Proposition 2.

$$(c_1 \mid s_1 \models A_1) \xRightarrow{\text{st}} (c_2 \mid s_2 \models A_2)$$

if, and only if,

$$(c_1 \wedge c_3 \mid s_1 \models A_1) \xRightarrow{\text{st}} (c_2 \wedge c_3 \mid s_2 \models A_2)$$

provided $FV(c_3) \cap SV(s_1, s_2) = \emptyset$ and $FV(c_3) \cap \alpha\text{st} = \emptyset$.

We use $SV(s)$ to denote the set of symbolic variables used in the definition of s (or, more generally, in a list of states s_1, s_2 as used in the proposition above). By construction, constraints of a configuration are only ever strengthened.

Proposition 3.

$$(c_1 \mid s_1 \models A_1) \xRightarrow{\text{st}} (c_2 \mid s_2 \models A_2)$$

if, and only if, for some c_3

$$(c_1 \mid s_1 \models A_1) \xRightarrow{\text{st}} (c_1 \wedge c_3 \mid s_2 \models A_2)$$

So, the constraint c_1 used in an initial configuration is preserved in any configuration reached from it.

The operational semantics is consistent with the denotational semantics. So, for example, we have the property below for traces of an external choice.

Proposition 4.

$$(c_1 \mid s_1 \models A_1 \sqcap A_2) \xRightarrow{\text{st}} (c_2 \mid s_2 \models A_3)$$

for some s_2 and A_3 if, and only if,

$$(c_1 \mid s_1 \models A_1) \xRightarrow{\text{st}} (c_2 \mid s_2 \models A_3) \quad \text{or} \quad (c_1 \mid s_1 \models A_2) \xRightarrow{\text{st}} (c_2 \mid s_2 \models A_3)$$

This is consistent with the denotational semantics of $A_1 \sqcap A_2$, which defines that its set of traces is the union of those of A_1 and A_2 . For sequential composition, we have the following.

Proposition 5.

$$(c_1 \mid s_1 \models A_1 ; A_2) \xRightarrow{\text{st}} (c_2 \mid s_2 \models A_3)$$

for some s_2 and A_3 if, and only if,

$$(c_1 \mid s_1 \models A_1) \xRightarrow{\text{st}} (c_2 \mid s_2 \models A_3)$$

or, there are st_1, c'_1, s'_1 , and st_2 such that $\text{st} = \text{st}_1 \hat{\ } \text{st}_2$,

$$(c_1 \mid s_1 \models A_1) \xRightarrow{\text{st}_1} (c'_1 \mid s'_1 \models \text{Skip}) \quad \text{and} \quad (c'_1 \mid s'_1 \models A_2) \xRightarrow{\text{st}_2} (c_2 \mid s_2 \models A_3)$$

The next property is concerned with the semantics of parallel actions. It relates the traces of a parallelism $A_1 \llbracket x_1 \mid cs \mid x_2 \rrbracket A_2$ of actions, as calculated by the operational semantics, with the result of applying the synchronisation operator to the traces st_1 and st_2 of the parallel actions A_1 and A_2 . This is needed in this work, where we consider parallel processes, because parallelism of processes is defined in terms of the parallelism of their main actions, with name sets defined by the states of the processes.

Proposition 6.

$$(c_1 \mid s_1 \models A_1 \llbracket x_1 \mid cs \mid x_2 \rrbracket A_2) \xRightarrow{st} (c_2 \mid s_2 \models A_3)$$

for some s_2 and A_3 if, and only if, there are $st_1, c'_1, s'_1, A'_1, st_2, c'_2, s'_2, A'_2$, and a , such that $\alpha c_1' \cap \alpha c_2' = \emptyset$, $\alpha st_1 \subseteq \alpha c_1'$, $\alpha st_2 \subseteq \alpha c_2'$, $\alpha st \leq a$,

$$(c_1 \mid s_1 \models A_1) \xRightarrow{st_1} (c_1' \mid s_1' \models A_1'), \quad (c_1 \mid s_1 \models A_2) \xRightarrow{st_2} (c_2' \mid s_2' \models A_2')$$

and

$$\mathcal{R}(st, c_2) \in \mathcal{N}^a (\mathcal{R} (\llbracket st_1 \llbracket cs, (c_1' \wedge c_2') \rrbracket st_2 \rrbracket))$$

We need to apply \mathcal{R} to all constrained symbolic traces: that obtained from the operational semantics of $A_1 \llbracket x_1 \mid cs \mid x_2 \rrbracket A_2$ and those obtained by the synchronisation of st_1 and st_2 . The renaming carried out by \mathcal{N}^a ensures that all symbolic traces are over the same alphabet a , and so, can be compared.

Finally, we have a property of the operational semantics of hiding.

Proposition 7.

$$(c_1 \mid s_1 \models A \setminus cs) \xRightarrow{st_1} (c_2 \mid s_2 \models A_2)$$

for some s_2 and A_2 if, and only if, there are st_2, c_3, s_3, A_3 , and a , such that $\alpha st_1 \leq a$,

$$(c_1 \mid s_1 \models A) \xRightarrow{st_2} (c_3 \mid s_3 \models A_3) \quad \text{and} \quad \mathcal{R}(st_1, c_2) = \mathcal{N}^a(\mathcal{R}(st_2, c_3) \upharpoonright_C cs)$$

We consider traces st_2 defined by the evaluation of A , and their projections to remove the channels in cs .

Appendix C. Some lemmas and their proofs

Lemma 1.

$$Scstraces_{cs}^{a, (a_1, a_2)}(P_1, P_2) = Scstraces_{cs}^{a, (a_3, a_4)}(P_1, P_2)$$

Proof. There is a bijection between $cstraces^{a_1}(P_1)$ and $cstraces^{a_3}(P_1)$ reflecting the fact that their symbolic traces characterise the same set of traces, namely those of P_1 , and the only difference between the symbolic traces in these sets is in their use of names of symbolic variables. The same comment applies to $cstraces^{a_2}(P_2)$ and $cstraces^{a_4}(P_2)$. We, therefore, can conclude that there is such a bijection between $ScstraceP_{cs}^{(a_1, a_2)}(P_1, P_2)$ and $Scstraces_{cs}^{(a_3, a_4)}(P_1, P_2)$. So, this lemma is a direct consequence of the definition of \mathcal{N}^a . \square

Appendix C.1. Properties of synchronisation

The constraints of all traces are the same.

Lemma 2.

$$\forall cst : st_1 \llbracket cs, c \rrbracket st_2 \bullet cst.2 = (st_1 \upharpoonright_E cs =_S st_2 \upharpoonright_E cs) \wedge c$$

Proof. By induction.

Case 1.

$$\begin{aligned} (st, c_1) \in \langle \rangle \llbracket cs, c_2 \rrbracket \langle \rangle \\ \Rightarrow c_1 = c_2 & \hspace{15em} [\text{definition of } (\langle \rangle \llbracket cs, c_2 \rrbracket \langle \rangle)] \\ \Leftrightarrow c_1 = ((\langle \rangle \upharpoonright_E cs) =_S (\langle \rangle \upharpoonright_E cs)) \wedge c_2 & \hspace{15em} [\text{definition of } \upharpoonright_E \text{ and } =_S] \end{aligned}$$

Case 2.

$$\begin{aligned}
& (\mathbf{st}, c_1) \in \langle \rangle \llbracket \mathbf{cs}, c_2 \rrbracket \langle \mathbf{d}.\beta \rangle \\
& \Rightarrow c_1 = c_2 \quad \text{[definition of } \langle \rangle \llbracket \mathbf{cs}, c_2 \rrbracket \langle \mathbf{d}.\beta \rangle] \\
& \Leftrightarrow c_1 = (((\langle \rangle \upharpoonright_E \mathbf{cs}) =_S (\langle \mathbf{d}.\beta \rangle \upharpoonright_E \mathbf{cs})) \wedge c_2 \quad \text{[[} \langle \mathbf{d}.\beta \rangle \upharpoonright_E \mathbf{cs} \rangle = \langle \rangle]
\end{aligned}$$

Case 3.: $d_1 \in \mathbf{cs}, d_2 \notin \mathbf{cs}$

$$\begin{aligned}
& (\mathbf{st}, c_1) \in (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2) \\
& \Rightarrow \exists \mathbf{st}_3, c_3 \bullet \\
& \quad (\mathbf{st}_3, c_3) \in (((\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket \mathbf{st}_2) \wedge \mathbf{st}_1 = \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_3 \wedge c_1 = c_3 \\
& \quad \quad \quad \text{[definition of } (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2)] \\
& \Leftrightarrow \exists \mathbf{st}_3 \bullet (\mathbf{st}_3, c_1) \in (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket \mathbf{st}_2 \wedge \mathbf{st}_1 = \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_3 \quad \text{[predicate calculus]} \\
& \Rightarrow c_1 = (((\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \upharpoonright_E \mathbf{cs}) =_S \mathbf{st}_2 \upharpoonright_E \mathbf{cs}) \wedge c_2 \quad \text{[induction hypothesis]} \\
& \Leftrightarrow c_1 = (((\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \upharpoonright_E \mathbf{cs}) =_S (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2) \upharpoonright_E \mathbf{cs}) \wedge c_2 \quad \text{[[} \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2 \upharpoonright_E \mathbf{cs} = \mathbf{st}_2 \upharpoonright_E \mathbf{cs}]]
\end{aligned}$$

Case 4.: $d \in \mathbf{cs}$

$$\begin{aligned}
& (\mathbf{st}, c_1) \in (\langle \mathbf{d}.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket (\langle \mathbf{d}.\beta \rangle \wedge \mathbf{st}_2) \\
& \Leftrightarrow \exists \mathbf{st}_3 \bullet (\mathbf{st}_3, c_1) \in \mathbf{st}_1 \llbracket \mathbf{cs}, (c_2 \wedge \alpha = \beta) \rrbracket \mathbf{st}_2 \wedge \mathbf{st}_1 = \langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_3 \\
& \quad \quad \quad \text{[definition of } (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2)] \\
& \Rightarrow c_1 = (\mathbf{st}_1 \upharpoonright_E \mathbf{cs} =_S \mathbf{st}_2 \upharpoonright_E \mathbf{cs}) \wedge (c_2 \wedge \alpha = \beta) \quad \text{[induction hypothesis]} \\
& \Leftrightarrow c_1 = (\langle \mathbf{d}.\alpha \rangle \wedge (\mathbf{st}_1 \upharpoonright_E \mathbf{cs}) =_S \langle \mathbf{d}.\beta \rangle \wedge (\mathbf{st}_2 \upharpoonright_E \mathbf{cs})) \wedge c_2 \quad \text{[definition of } =_S] \\
& \Leftrightarrow c_1 = (((\langle \mathbf{d}.\alpha \rangle \wedge \mathbf{st}_1) \upharpoonright_E \mathbf{cs}) =_S (\langle \mathbf{d}.\beta \rangle \wedge \mathbf{st}_2) \upharpoonright_E \mathbf{cs}) \wedge c_2 \\
& \quad \quad \quad \text{[[} (\langle \mathbf{d}.\alpha \rangle \wedge \mathbf{st}_1) \upharpoonright_E \mathbf{cs} = \langle \mathbf{d}.\alpha \rangle \wedge (\mathbf{st}_1 \upharpoonright_E \mathbf{cs}) \text{ and } (\langle \mathbf{d}.\beta \rangle \wedge \mathbf{st}_2) \upharpoonright_E \mathbf{cs} = \langle \mathbf{d}.\beta \rangle \wedge (\mathbf{st}_2 \upharpoonright_E \mathbf{cs})]]
\end{aligned}$$

Case 5.: $d_1 \notin \mathbf{cs}, d_2 \notin \mathbf{cs}$

$$\begin{aligned}
& (\mathbf{st}, c_1) \in (((\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2)) \\
& \Rightarrow \exists \mathbf{st}_3 \bullet (\mathbf{st}_3, c_1) \in (((\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket \mathbf{st}_2) \vee \\
& \quad (\mathbf{st}_3, c_1) \in (\mathbf{st}_1 \llbracket \mathbf{cs}, c_2 \rrbracket (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2)) \\
& \quad \quad \quad \text{[definition of } (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \llbracket \mathbf{cs}, c_2 \rrbracket (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2)] \\
& \Rightarrow c_1 = (((\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \upharpoonright_E \mathbf{cs}) =_S \mathbf{st}_2 \upharpoonright_E \mathbf{cs}) \wedge c_2 \vee \quad \text{[induction hypothesis]} \\
& \quad c_1 = (\mathbf{st}_1 \upharpoonright_E \mathbf{cs} =_S (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2) \upharpoonright_E \mathbf{cs}) \wedge c_2 \\
& \Rightarrow c_1 = (((\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \upharpoonright_E \mathbf{cs}) =_S (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2) \upharpoonright_E \mathbf{cs}) \wedge c_2 \\
& \quad \quad \quad \text{[[} (\langle \mathbf{d}_1.\alpha \rangle \wedge \mathbf{st}_1) \upharpoonright_E \mathbf{cs} = (\mathbf{st}_1 \upharpoonright_E \mathbf{cs}) \text{ and } (\langle \mathbf{d}_2.\beta \rangle \wedge \mathbf{st}_2) \upharpoonright_E \mathbf{cs} = (\mathbf{st}_2 \upharpoonright_E \mathbf{cs})]]
\end{aligned}$$

□

The projection of the synchronisation traces in $\mathbf{st}_1 \llbracket \mathbf{cs}, c \rrbracket \mathbf{st}_2$ to their symbolic traces does not change if we quantify one of the variables (potentially) free in c .

Lemma 3.

$$(\mathbf{st}_1 \llbracket \mathbf{cs}, c \rrbracket \mathbf{st}_2) .1 = (\mathbf{st}_1 \llbracket \mathbf{cs}, (\exists x \bullet c) \rrbracket \mathbf{st}_2) .1$$

Proof. By induction.

Case 1(a). : provided c

$$\begin{aligned}
& \langle\langle \llbracket \text{cs}, c \rrbracket \rangle\rangle .1 \\
&= \langle\{ \langle\langle \rangle, c \rangle \} \rangle .1 && \text{[synchronisation operator]} \\
&= \{ \langle\langle \rangle \} && \text{[relational image]} \\
&= \langle\{ \langle\langle \rangle, (\exists x \bullet c) \rangle \} \rangle .1 && \text{[relational image]} \\
&= \langle\langle \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \rangle\rangle .1 && \text{[synchronisation operator]}
\end{aligned}$$

Case 1(b). : otherwise

$$\begin{aligned}
& \langle\langle \llbracket \text{cs}, c \rrbracket \rangle\rangle .1 \\
&= \emptyset && \text{[synchronisation operator and relational image]} \\
&= \langle\langle \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \rangle\rangle .1 && \text{[(}\exists x \bullet c\text{) does not hold]}
\end{aligned}$$

Case 2(a). : $\langle\langle \llbracket \text{cs}, c_2 \rrbracket \langle d.\beta \rangle \rangle\rangle$, provided c

$$\begin{aligned}
& \langle\langle \llbracket \text{cs}, c \rrbracket \langle d.\beta \rangle \rangle\rangle .1 \\
&= \langle\{ \langle\langle d.\beta \rangle, c \rangle \} \rangle .1 && \text{[synchronisation operator]} \\
&= \{ \langle d.\beta \rangle \} && \text{[relational image]} \\
&= \langle\{ \langle\langle d.\beta \rangle, (\exists x \bullet c) \rangle \} \rangle .1 && \text{[relational image]} \\
&= \langle\langle \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \langle d.\beta \rangle \rangle\rangle .1 && \text{[synchronisation operator]}
\end{aligned}$$

Case 2(a). : $\langle\langle \llbracket \text{cs}, c \rrbracket \langle d.\beta \rangle \rangle\rangle$, provided c does not hold
Similar to Case 1(b).

Case 3: $d_1 \in \text{cs}, d_2 \notin \text{cs}$

$$\begin{aligned}
& \langle\langle \langle d_1.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, c \rrbracket \langle \langle d_2.\beta \rangle \wedge \text{st}_2 \rangle \rangle\rangle .1 \\
&= \langle\{ \text{cst} : \langle \langle d_1.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, c \rrbracket \text{st}_2 \bullet \langle \langle d_2.\beta \rangle \wedge \text{cst}.1, \text{cst}.2 \rangle \} \} \rangle .1 && \text{[synchronisation operator]} \\
&= \langle\{ \text{cst}_1 : \langle \langle d_1.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \text{st}_2; \text{cst}_2 : \langle \langle d_1.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, c \rrbracket \text{st}_2 \bullet \langle \langle d_2.\beta \rangle \wedge \text{cst}_1.1, \text{cst}_2.2 \rangle \} \} \rangle .1 && \text{[induction hypothesis]} \\
&= \{ \text{cst} : \langle \langle d_1.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \text{st}_2 \bullet \langle d_2.\beta \rangle \wedge \text{cst}.1 \} && \text{[relational image]} \\
&= \langle\{ \text{cst} : \langle \langle d_1.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \text{st}_2 \bullet \langle \langle d_2.\beta \rangle \wedge \text{cst}.1, \text{cst}.2 \rangle \} \} \rangle .1 && \text{[relational image]} \\
&= \langle\langle \langle d_1.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \langle \langle d_2.\beta \rangle \wedge \text{st}_2 \rangle \rangle\rangle .1 && \text{[synchronisation operator]}
\end{aligned}$$

Case 4: $d \in \text{cs}$

$$\begin{aligned}
& \langle\langle \langle d.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, c \rrbracket \langle \langle d.\beta \rangle \wedge \text{st}_2 \rangle \rangle\rangle .1 \\
&= \langle\{ \text{cst} : \text{st}_1 \llbracket \text{cs}, c \rrbracket \text{st}_2 \bullet \langle \langle d.\beta \rangle \wedge \text{cst}.1, \text{cst}.2 \rangle \} \} \rangle .1 && \text{[synchronisation operator]} \\
&= \langle\{ \text{cst}_1 : \text{st}_1 \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \text{st}_2; \text{cst}_2 : \text{st}_1 \llbracket \text{cs}, c \rrbracket \text{st}_2 \bullet \langle \langle d.\beta \rangle \wedge \text{cst}_1.1, \text{cst}_2.2 \rangle \} \} \rangle .1 && \text{[induction hypothesis]} \\
&= \{ \text{cst} : \text{st}_1 \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \text{st}_2 \bullet \langle d.\beta \rangle \wedge \text{cst}.1 \} && \text{[relational image]} \\
&= \langle\{ \text{cst} : \text{st}_1 \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \text{st}_2 \bullet \langle \langle d.\beta \rangle \wedge \text{cst}.1, \text{cst}.2 \rangle \} \} \rangle .1 && \text{[relational image]} \\
&= \langle\langle \langle d.\alpha \rangle \wedge \text{st}_1 \llbracket \text{cs}, (\exists x \bullet c) \rrbracket \langle \langle d.\beta \rangle \wedge \text{st}_2 \rangle \rangle\rangle .1 && \text{[synchronisation operator]}
\end{aligned}$$

Case 5: $(\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, c \rrbracket (\langle d_2.\beta \rangle \wedge st_2)$, with $d_1 \neq d_2, d_1 \in cs, d_2 \in cs$
 Similar to Case 1(b).

Case 6: $d_1 \notin cs, d_2 \notin cs$

$$\begin{aligned}
 & \llbracket (\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, c \rrbracket (\langle d_2.\beta \rangle \wedge st_2) \rrbracket .1 \\
 = & \llbracket \{ cst : (\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, c \rrbracket st_2 \bullet (\langle d_2.\beta \rangle \wedge cst.1, cst.2) \} \cup \quad [\text{synchronisation operator}] \\
 & \{ cst : st_1 \llbracket cs, c \rrbracket (\langle d_2.\beta \rangle \wedge st_2) \bullet (\langle d_1.\alpha \rangle \wedge cst.1, cst.2) \} \rrbracket .1 \\
 = & \llbracket \{ cst : (\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, c \rrbracket st_2 \bullet (\langle d_2.\beta \rangle \wedge cst.1, cst.2) \} \rrbracket .1 \cup \quad [\text{property of relational image}] \\
 & \llbracket \{ cst : st_1 \llbracket cs, c \rrbracket (\langle d_2.\beta \rangle \wedge st_2) \bullet (\langle d_1.\alpha \rangle \wedge cst.1, cst.2) \} \rrbracket .1 \\
 = & \llbracket \{ cst : (\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, (\exists x \bullet c) \rrbracket st_2 \bullet (\langle d_2.\beta \rangle \wedge cst.1, cst.2) \} \rrbracket .1 \cup \quad [\text{induction hypothesis}] \\
 & \llbracket \{ cst : st_1 \llbracket cs, (\exists x \bullet c) \rrbracket (\langle d_2.\beta \rangle \wedge st_2) \bullet (\langle d_1.\alpha \rangle \wedge cst.1, cst.2) \} \rrbracket .1 \\
 = & \llbracket \{ cst : (\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, (\exists x \bullet c) \rrbracket st_2 \bullet (\langle d_2.\beta \rangle \wedge cst.1, cst.2) \} \cup \\
 & \{ cst : st_1 \llbracket cs, (\exists x \bullet c) \rrbracket (\langle d_2.\beta \rangle \wedge st_2) \bullet (\langle d_1.\alpha \rangle \wedge cst.1, cst.2) \} \rrbracket .1 \\
 & \quad \quad \quad [\text{property of relational image}] \\
 = & \llbracket (\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, (\exists x \bullet c) \rrbracket (\langle d_2.\beta \rangle \wedge st_2) \rrbracket .1 \quad [\text{synchronisation operator}]
 \end{aligned}$$

□

In the lemma below, we have a simple property of \mathcal{R} when used in conjunction with the synchronisation operator. Basically, we can, in advance, consider a weaker constraint that already quantifies away variables that represent internal values: those that are not used in either of the traces st_1 and st_2 synchronised.

Lemma 4.

$$\mathcal{R} (\llbracket st_1 \llbracket cs, c \rrbracket st_2 \rrbracket) = \mathcal{R} (\llbracket st_1 \llbracket cs, (\exists (\alpha c \setminus \alpha(st_1, st_2)) \bullet c) \rrbracket st_2 \rrbracket)$$

Proof.

$$\begin{aligned}
 & \mathcal{R} (\llbracket st_1 \llbracket cs, c \rrbracket st_2 \rrbracket) \\
 = & \{ cst : st_1 \llbracket cs, c \rrbracket st_2 \bullet \mathcal{R}(cst.1, cst.2) \} \quad [\text{relational image}] \\
 = & \{ cst : st_1 \llbracket cs, c \rrbracket st_2 \bullet \mathcal{R}(cst.1, (st_1 \upharpoonright_E cs =_S st_2 \upharpoonright_E cs) \wedge c) \} \quad [\text{Lemma 2}] \\
 = & \{ cst : st_1 \llbracket cs, c \rrbracket st_2 \bullet \mathcal{R}(cst.1, \exists x_1 \bullet (st_1 \upharpoonright_E cs =_S st_2 \upharpoonright_E cs) \wedge c) \} \\
 & \quad \quad \quad [\text{definition of } \mathcal{R}, \text{ with } x_1 = \alpha(st_1 \upharpoonright_E cs, st_2 \upharpoonright_E cs, c) \setminus \alpha cst.1] \\
 = & \{ cst : st_1 \llbracket cs, c \rrbracket st_2 \bullet (cst.1, \exists x_2 \bullet (st_1 \upharpoonright_E cs =_S st_2 \upharpoonright_E cs) \wedge c) \} \\
 & \quad \quad \quad [\text{property of sets, with } x_2 = (\alpha(st_1 \upharpoonright_E cs, st_2 \upharpoonright_E cs), \alpha c \cap \alpha(st_1, st_2), \alpha c \setminus \alpha(st_1, st_2)) \setminus \alpha cst.1] \\
 = & \{ cst : st_1 \llbracket cs, c \rrbracket st_2 \bullet (cst.1, \exists x_3 \bullet (st_1 \upharpoonright_E cs =_S st_2 \upharpoonright_E cs) \wedge c) \} \\
 & \quad \quad \quad [\alpha cst.1 \subseteq \alpha(st_1, st_2), \text{ with } x_3 = (\alpha(st_1 \upharpoonright_E cs, st_2 \upharpoonright_E cs), \alpha c \cap \alpha(st_1, st_2)) \setminus \alpha cst.1, \alpha c \setminus \alpha(st_1, st_2)] \\
 = & \{ cst : st_1 \llbracket cs, c \rrbracket st_2 \bullet (cst.1, \exists x_4 \bullet (st_1 \upharpoonright_E cs =_S st_2 \upharpoonright_E cs) \wedge \exists (\alpha c \setminus \alpha(st_1, st_2)) \bullet c) \} \\
 & \quad \quad \quad [\text{predicate calculus, with } x_4 = (\alpha(st_1 \upharpoonright_E cs, st_2 \upharpoonright_E cs), \alpha c \cap \alpha(st_1, st_2)) \setminus \alpha cst.1]
 \end{aligned}$$

$$\begin{aligned}
&= \{ \text{cst} : \text{st}_1 \llbracket \text{cs}, (\exists(\alpha c \setminus \alpha(\text{st}_1, \text{st}_2)) \bullet c) \rrbracket \text{st}_2 \bullet && \text{[Lemma 3]} \\
&\quad (\text{cst}.1, \exists x_4 \bullet (\text{st}_1 \upharpoonright_E \text{cs} =_S \text{st}_2 \upharpoonright_E \text{cs}) \wedge \exists(\alpha c \setminus \alpha(\text{st}_1, \text{st}_2)) \bullet c) \} \\
&= \{ \text{cst} : \text{st}_1 \llbracket \text{cs}, (\exists(\alpha c \setminus \alpha(\text{st}_1, \text{st}_2)) \bullet c) \rrbracket \text{st}_2 \bullet && \text{[definition of } \mathcal{R} \text{]} \\
&\quad \mathcal{R}(\text{cst}.1, (\text{st}_1 \upharpoonright_E \text{cs} =_S \text{st}_2 \upharpoonright_E \text{cs}) \wedge \exists(\alpha c \setminus \alpha(\text{st}_1, \text{st}_2)) \bullet c) \} \\
&= \{ \text{cst} : \text{st}_1 \llbracket \text{cs}, (\exists(\alpha c \setminus \alpha(\text{st}_1, \text{st}_2)) \bullet c) \rrbracket \text{st}_2 \bullet \mathcal{R}(\text{cst}) \} && \text{[Lemma 2]} \\
&= \mathcal{R}(\llbracket \text{st}_1 \llbracket \text{cs}, (\exists(\alpha c \setminus \alpha(\text{st}_1, \text{st}_2)) \bullet c) \rrbracket \text{st}_2) && \text{[relational image]}
\end{aligned}$$

□

The following lemma establishes that the set of traces $\text{st}_1 \llbracket \text{cs}, c \rrbracket \text{st}_2$ is not empty exactly when c is satisfiable, even in the presence of the synchronisation requirements imposed by st_1 and st_2 .

Lemma 5.

$$(\text{st}_1 \llbracket \text{cs}, c \rrbracket \text{st}_2) \neq \emptyset \Leftrightarrow \exists a_1, a_2 \bullet (\text{st}_1 \upharpoonright_E \text{cs}) =_S (\text{st}_2 \upharpoonright_E \text{cs}) \wedge c$$

where $\alpha \text{st}_1 \leq a_1$ and $\alpha \text{st}_2 \leq a_2$.

Proof. By induction.

Case 1.

$$\begin{aligned}
&\langle \rangle \llbracket \text{cs}, c \rrbracket \langle \rangle \neq \emptyset \\
&\Leftrightarrow c && \text{[definition of } \langle \rangle \llbracket \text{cs}, c \rrbracket \langle \rangle \text{]} \\
&\Leftrightarrow \exists a_1, a_2 \bullet \langle \rangle \upharpoonright_E \text{cs} =_S \langle \rangle \upharpoonright_E \text{cs} \wedge c && [\langle \rangle \upharpoonright_E \text{cs} = \langle \rangle \text{ and } \langle \rangle =_S \langle \rangle = \text{true}]
\end{aligned}$$

Case 2(a). : $d \notin \text{cs}$

$$\begin{aligned}
&\langle \rangle \llbracket \text{cs}, c \rrbracket \langle d.\beta \rangle \neq \emptyset \\
&\Leftrightarrow c && \text{[definition of } \langle \rangle \llbracket \text{cs}, c \rrbracket \langle d.\beta \rangle \text{]} \\
&\Leftrightarrow \exists a_1, a_2 \bullet \langle \rangle \upharpoonright_E \text{cs} =_S \langle d.\beta \rangle \upharpoonright_E \text{cs} \wedge c && [\langle d.\beta \rangle \upharpoonright_E \text{cs} = \langle \rangle]
\end{aligned}$$

Case 2(b). : $d \in \text{cs}$

$$\begin{aligned}
&\langle \rangle \llbracket \text{cs}, c \rrbracket \langle d.\beta \rangle \neq \emptyset \\
&\Leftrightarrow \text{false} && \text{[definition of } \langle \rangle \llbracket \text{cs}, c \rrbracket \langle d.\beta \rangle \text{]} \\
&\Leftrightarrow \exists a_1, a_2 \bullet \langle \rangle \upharpoonright_E \text{cs} =_S \langle d.\beta \rangle \upharpoonright_E \text{cs} \wedge c && [\langle d.\beta \rangle \upharpoonright_E \text{cs} = \langle d.\beta \rangle \text{ and } \langle \rangle =_S \langle d.\beta \rangle = \text{false}]
\end{aligned}$$

Case 3: $d_1 \in \text{cs}, d_2 \notin \text{cs}$

$$\begin{aligned}
&((\langle d_1.\alpha \rangle \wedge \text{st}_1) \llbracket \text{cs}, c \rrbracket ((\langle d_2.\beta \rangle \wedge \text{st}_2)) \neq \emptyset \\
&\Leftrightarrow ((\langle d_1.\alpha \rangle \wedge \text{st}_1) \llbracket \text{cs}, c \rrbracket \text{st}_2) \neq \emptyset && \text{[definition of } (\langle d_1.\alpha \rangle \wedge \text{st}_1) \llbracket \text{cs}, c \rrbracket ((\langle d_2.\beta \rangle \wedge \text{st}_2)) \text{]} \\
&\Leftrightarrow \exists a_1, a_2 \bullet (((\langle d_1.\alpha \rangle \wedge \text{st}_1) \upharpoonright_E \text{cs}) =_S (\text{st}_2 \upharpoonright_E \text{cs}) \wedge c && \text{[induction hypothesis]} \\
&\Leftrightarrow \exists a_1, a_2 \bullet (((\langle d_1.\alpha \rangle \wedge \text{st}_1) \upharpoonright_E \text{cs}) =_S (((\langle d_2.\beta \rangle \wedge \text{st}_2) \upharpoonright_E \text{cs}) \wedge c && [((\langle d_2.\beta \rangle \wedge \text{st}_2) \upharpoonright_E \text{cs}) = (\text{st}_2 \upharpoonright_E \text{cs})]
\end{aligned}$$

Case 4: $d \in \text{cs}$

$$\begin{aligned}
&((\langle d.\alpha \rangle \wedge \text{st}_1) \llbracket \text{cs}, c \rrbracket ((\langle d.\beta \rangle \wedge \text{st}_2)) \neq \emptyset \\
&\Leftrightarrow (\text{st}_1 \llbracket \text{cs}, (c \wedge \alpha = \beta) \rrbracket \text{st}_2) \neq \emptyset && \text{[definition of } (\langle d.\alpha \rangle \wedge \text{st}_1) \llbracket \text{cs}, c \rrbracket ((\langle d.\beta \rangle \wedge \text{st}_2)) \text{]} \\
&\Leftrightarrow \exists a_1, a_2 \bullet (\text{st}_1 \upharpoonright_E \text{cs}) =_S (\text{st}_2 \upharpoonright_E \text{cs}) \wedge c \wedge \alpha = \beta && \text{[induction hypothesis]} \\
&\Leftrightarrow \exists a_1, a_2 \bullet (\langle d.\alpha \rangle \wedge (\text{st}_1 \upharpoonright_E \text{cs})) =_S (\langle d.\beta \rangle \wedge (\text{st}_2 \upharpoonright_E \text{cs})) \wedge c && \text{[definition of } =_S \text{]} \\
&\Leftrightarrow \exists a_1, a_2 \bullet (((\langle d.\alpha \rangle \wedge \text{st}_1) \upharpoonright_E \text{cs}) =_S (((\langle d.\beta \rangle \wedge \text{st}_2) \upharpoonright_E \text{cs}) \wedge c && \text{[definition of } \upharpoonright_E \text{]}
\end{aligned}$$

Case 5.: $d_1 \neq d_2, d_1 \in cs, d_2 \in cs$

$$(((d_1.\alpha) \wedge st_1) \llbracket cs, c \rrbracket ((d_2.\beta) \wedge st_2)) \neq \emptyset$$

$$\Leftrightarrow \text{false} \quad [\text{definition of } ((d_1.\alpha) \wedge st_1) \llbracket cs, c \rrbracket ((d_2.\beta) \wedge st_2)]$$

$$\Leftrightarrow \exists a_1, a_2 \bullet (((d_1.\alpha) \wedge st_1) \downarrow_E cs) =_S (((d_2.\beta) \wedge st_2) \downarrow_E cs) \wedge c$$

$$[\text{if } (((d_1.\alpha) \wedge st_1) \downarrow_E cs) =_S (((d_2.\beta) \wedge st_2) \downarrow_E cs) \text{ is false}]$$

Case 6.: $d_1 \notin cs, d_2 \notin cs$

$$(((d_1.\alpha) \wedge st_1) \llbracket cs, c \rrbracket ((d_2.\beta) \wedge st_2)) \neq \emptyset$$

$$\Leftrightarrow (((d_1.\alpha) \wedge st_1) \llbracket cs, c \rrbracket st_2) \neq \emptyset \vee (st_1 \llbracket cs, c \rrbracket ((d_2.\beta) \wedge st_2)) \neq \emptyset$$

$$[\text{definition of } ((d_1.\alpha) \wedge st_1) \llbracket cs, c \rrbracket ((d_2.\beta) \wedge st_2)]$$

$$\Leftrightarrow (\exists a_1, a_2 \bullet (((d_1.\alpha) \wedge st_1) \downarrow_E cs) =_S (st_2 \downarrow_E cs) \wedge c) \vee$$

[induction hypothesis]

$$(\exists a_1, a_2 \bullet (st_1 \downarrow_E cs) =_S (((d_2.\beta) \wedge st_2) \downarrow_E cs) \wedge c)$$

$$\Leftrightarrow \exists a_1, a_2 \bullet (((d_1.\alpha) \wedge st_1) \downarrow_E cs) =_S (((d_2.\beta) \wedge st_2) \downarrow_E cs) \wedge c$$

$$[\text{if } ((d_1.\alpha) \wedge st_1) \downarrow_E cs = (st_1 \downarrow_E cs) \text{ and } ((d_2.\beta) \wedge st_2) \downarrow_E cs = (st_2 \downarrow_E cs)]$$

□

With the next lemma, we establish that, if the projection of both traces st_1 and st_2 to synchronisation channels is nonempty, then the projection of all synchronisation traces to these channels is nonempty.

Lemma 6.

$$\forall st, c_1 \bullet (st, c_1) \in (st_1 \llbracket cs, c_2 \rrbracket st_2) \Rightarrow st \downarrow_C cs \neq \langle \rangle \wedge st_1 \downarrow_C cs \neq \langle \rangle \wedge st_2 \downarrow_C cs \neq \langle \rangle$$

provided $(st_1 \llbracket cs, c_2 \rrbracket st_2) \neq \emptyset$.

Proof. By induction.

Case 1.

$$\forall st, c_1 \bullet (st, c_1) \in (\langle \rangle \llbracket cs, c_2 \rrbracket \langle \rangle) \Rightarrow st \downarrow_C cs \neq \langle \rangle$$

$$\Leftrightarrow \text{false}$$

[definitions of $(\langle \rangle \llbracket cs, c_2 \rrbracket \langle \rangle)$ and \downarrow_C]

$$\Leftrightarrow \langle \rangle \downarrow_C cs \neq \langle \rangle \wedge \langle \rangle \downarrow_C cs \neq \langle \rangle$$

$$[\langle \rangle \downarrow_C cs = \langle \rangle]$$

Case 2. : $d \notin cs$

$$\forall st, c_1 \bullet (st, c_1) \in (\langle \rangle \llbracket cs, c_2 \rrbracket \langle d.\beta \rangle) \Rightarrow st \downarrow_C cs \neq \langle \rangle$$

$$\Leftrightarrow \text{false}$$

[definition of $(\langle \rangle \llbracket cs, c_2 \rrbracket \langle d.\beta \rangle)$ and $\langle d.\beta \rangle \downarrow_C cs = \langle \rangle$, if $d \notin cs$]

$$\Leftrightarrow \langle \rangle \downarrow_C cs \neq \langle \rangle \wedge \langle d.\beta \rangle \downarrow_C cs \neq \langle \rangle$$

$$[\langle \rangle \downarrow_C cs = \langle \rangle \text{ (and } \langle d.\beta \rangle \downarrow_C cs = \langle \rangle)]$$

Case 3.: $d_1 \in cs, d_2 \notin cs$

$$\forall st, c_1 \bullet (st, c_1) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d_2.\beta) \wedge st_2)) \Rightarrow st \downarrow_C cs \neq \langle \rangle$$

$$\Leftrightarrow \forall st, c_1 \bullet (st, c_1) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket st_2) \Rightarrow ((d_2.\beta) \wedge st) \downarrow_C cs \neq \langle \rangle$$

[definition of $((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d_2.\beta) \wedge st_2)$]

$$\Leftrightarrow \forall st, c_1 \bullet (st, c_1) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket st_2) \Rightarrow st \downarrow_C cs \neq \langle \rangle$$

[definition of \downarrow_C]

$$\Leftrightarrow ((d_1.\alpha) \wedge st_1) \downarrow_C cs \neq \langle \rangle \wedge st_2 \downarrow_C cs \neq \langle \rangle$$

[induction hypothesis]

$$\Leftrightarrow ((d_1.\alpha) \wedge st_1) \downarrow_C cs \neq \langle \rangle \wedge ((d_2.\alpha) \wedge st_2) \downarrow_C cs \neq \langle \rangle$$

[definition of \downarrow_C]

Case 4: $d \in cs$

$$\begin{aligned}
& \forall st, c_1 \bullet (st, c_1) \in (((d.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d.\beta) \wedge st_2)) \Rightarrow st \downarrow_C cs \neq \langle \rangle \\
& \Leftrightarrow \forall st, c_1 \bullet (st, c_1) \in (st_1 \llbracket cs, (c_2 \wedge \alpha = \beta) \rrbracket st_2) \Rightarrow ((d.\beta) \wedge st) \downarrow_C cs \neq \langle \rangle \\
& \hspace{15em} [\text{definition of } ((d.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d.\beta) \wedge st_2)] \\
& \Leftrightarrow \text{true} \hspace{15em} [((d.\beta) \wedge st) \downarrow_C cs \neq \langle \rangle \text{ by definition of } \downarrow_C] \\
& \Leftrightarrow ((d.\alpha) \wedge st_1) \downarrow_C cs \neq \langle \rangle \wedge ((d.\beta) \wedge st_2) \downarrow_C cs \neq \langle \rangle \hspace{10em} [\text{definition of } \downarrow_C]
\end{aligned}$$

Case 5: $d_1 \notin cs, d_2 \notin cs$

$$\begin{aligned}
& \forall st, c_1 \bullet (st, c_1) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d_2.\beta) \wedge st_2)) \Rightarrow st \downarrow_C cs \neq \langle \rangle \\
& \Leftrightarrow (\forall st, c_1 \bullet (st, c_1) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket st_2) \Rightarrow ((d_2.\beta) \wedge st) \downarrow_C cs \neq \langle \rangle) \wedge \\
& \quad (\forall st, c_1 \bullet (st, c_1) \in (st_1 \llbracket cs, c_2 \rrbracket ((d_2.\beta) \wedge st_2)) \Rightarrow ((d_1.\alpha) \wedge st) \downarrow_C cs \neq \langle \rangle) \\
& \hspace{15em} [\text{definition of } (((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d_2.\beta) \wedge st_2))] \\
& \Leftrightarrow ((d_1.\alpha) \wedge st_1) \downarrow_C cs \neq \langle \rangle \wedge st_2 \downarrow_C cs \neq \langle \rangle \wedge \hspace{10em} [\text{induction hypothesis}] \\
& \quad st_1 \downarrow_C cs \neq \langle \rangle \wedge ((d_2.\beta) \wedge st_2) \downarrow_C cs \neq \langle \rangle \\
& \Leftrightarrow ((d_1.\alpha) \wedge st_1) \downarrow_C cs \neq \langle \rangle \wedge ((d_2.\beta) \wedge st_2) \downarrow_C cs \neq \langle \rangle \hspace{10em} [\text{definition of } \downarrow_C]
\end{aligned}$$

□

The following lemmas lift to the synchronisation operator for constrained symbolic traces results similar to some previously established for the operator for symbolic traces.

Lemma 7.

$$\forall cst \bullet cst \in \mathcal{N}^a (\llbracket cst_1 \rrbracket cs \rrbracket cst_2) \Rightarrow cst \downarrow_C cs \neq \langle \rangle \Leftrightarrow cst_1 \downarrow_C cs \neq \langle \rangle \wedge cst_2 \downarrow_C cs \neq \langle \rangle$$

where $\alpha st_1 \leq a_1$ and $\alpha st_2 \leq a_2$, and provided $\mathcal{N}^a (\llbracket cst_1 \rrbracket cs \rrbracket cst_2) \neq \emptyset$.

Proof.

$$\begin{aligned}
& \forall cst \bullet cst \in \mathcal{N}^a (\llbracket (st_1, c) \rrbracket cs \rrbracket (st_2, c)) \Rightarrow cst \downarrow_C cs \neq \langle \rangle \\
& \Leftrightarrow \forall st, c \bullet (st, c) \in ((st_1, c_1) \llbracket cs \rrbracket (st_2, c_2)) \Rightarrow (st, c) \downarrow_C cs \neq \langle \rangle \hspace{10em} [\text{definition of } \mathcal{N}] \\
& \Leftrightarrow \forall st, c \bullet (st, c) \in (st_1 \llbracket cs, (c_1 \wedge c_2) \rrbracket st_2) \Rightarrow (st, c) \downarrow_C cs \neq \langle \rangle \\
& \hspace{15em} [\text{definition of synchronising operator}] \\
& \Leftrightarrow \forall st, c \bullet (st, c) \in (st_1 \llbracket cs, (c_1 \wedge c_2) \rrbracket st_2) \Rightarrow st \downarrow_C cs \neq \langle \rangle \\
& \hspace{15em} [\text{definition of } \downarrow_C \text{ for constrained symbolic traces}] \\
& \Leftrightarrow st_1 \downarrow_C cs \neq \langle \rangle \wedge st_2 \downarrow_C cs \neq \langle \rangle \hspace{10em} [\text{Lemma 6 and proviso}] \\
& \Leftrightarrow (st_1, c_1) \downarrow_C cs \neq \langle \rangle \wedge (st_2, c_2) \downarrow_C cs \neq \langle \rangle \hspace{10em} [\text{definition of } \downarrow_C \text{ for constrained symbolic traces}]
\end{aligned}$$

□

Lemma 8.

$$\begin{aligned}
& \exists cst \bullet cst \in \mathcal{N}^a (\llbracket cst_1 \rrbracket cs \rrbracket cst_2) \wedge cst \downarrow_C cs \neq \langle \rangle \\
& \Rightarrow \\
& cst_1 \downarrow_C cs \neq \langle \rangle \wedge cst_2 \downarrow_C cs \neq \langle \rangle \wedge \text{satisfiable}_{cs}(cst_1, cst_2)
\end{aligned}$$

Proof.

$$\begin{aligned}
& \exists \text{cst} \bullet \text{cst} \in \mathcal{N}^a (\langle \text{st}_1, \text{c}_1 \rangle \parallel \text{cs} \parallel \langle \text{st}_2, \text{c}_2 \rangle) \wedge \text{cst} \downarrow_{\text{C}} \text{cs} \neq \langle \rangle \\
& \Leftrightarrow \exists \text{st}, \text{c} \bullet (\text{st}, \text{c}) \in (\langle \text{st}_1, \text{c}_1 \rangle \parallel \text{cs} \parallel \langle \text{st}_2, \text{c}_2 \rangle) \wedge (\text{st}, \text{c}) \downarrow_{\text{C}} \text{cs} \neq \langle \rangle \quad [\text{definition of } \mathcal{N}] \\
& \Leftrightarrow \exists \text{st}, \text{c} \bullet (\text{st}, \text{c}) \in (\text{st}_1 \parallel \text{cs}, (\text{c}_1 \wedge \text{c}_2) \parallel \text{st}_2) \wedge (\text{st}, \text{c}) \downarrow_{\text{C}} \text{cs} \neq \langle \rangle \quad [\text{definition of synchronisation operator}] \\
& \Leftrightarrow \exists \text{st}, \text{c} \bullet (\text{st}, \text{c}) \in (\text{st}_1 \parallel \text{cs}, (\text{c}_1 \wedge \text{c}_2) \parallel \text{st}_2) \wedge \text{st} \downarrow_{\text{C}} \text{cs} \neq \langle \rangle \\
& \quad \quad \quad [\text{definition of } \downarrow_{\text{C}} \text{ for constrained symbolic traces}] \\
& \Rightarrow \text{st}_1 \downarrow_{\text{C}} \text{cs} \neq \langle \rangle \wedge \text{st}_2 \downarrow_{\text{C}} \text{cs} \wedge \text{satisfiable}_{\text{cs}}(\langle \text{st}_1, \text{c}_1 \rangle, \langle \text{st}_2, \text{c}_2 \rangle) \quad [\text{Lemma 11}] \\
& \Leftrightarrow (\text{st}_1, \text{c}_1) \downarrow_{\text{C}} \text{cs} \neq \langle \rangle \wedge (\text{st}_2, \text{c}_2) \downarrow_{\text{C}} \text{cs} \wedge \text{satisfiable}_{\text{cs}}(\langle \text{st}_1, \text{c}_1 \rangle, \langle \text{st}_2, \text{c}_2 \rangle) \\
& \quad \quad \quad [\text{definition of } \downarrow_{\text{C}} \text{ for constrained symbolic traces}]
\end{aligned}$$

□

The lemma below establishes that the traces built from the (local) traces of P_1 and P_2 , are all those of the parallelism between P_1 and P_2 .

Lemma 9. *For every pair of disjoint alphabets \mathbf{a}_1 and \mathbf{a}_2*

$$\text{cstraces}^a(P_1 \parallel \text{cs} \parallel P_2) = \bigcup \{ \text{cst}_1 : \text{cstraces}^{a_1}(P_1); \text{cst}_2 : \text{cstraces}^{a_2}(P_2) \bullet \mathcal{N}^a (\langle \text{cst}_1 \parallel \text{cs} \parallel \text{cst}_2 \rangle) \}$$

Proof. In giving the semantics of $P_1 \parallel \text{cs} \parallel P_2$, we observe that we can assume, without loss of generality, that

$$P_1 = \text{begin state } [x_1 : T_1] \bullet A_1 \text{ end and } P_2 = \text{begin state } [x_2 : T_2] \bullet A_2 \text{ end}$$

with $x_1 \cap x_2 = \emptyset$, where x_1 and x_2 are the lists of state components of P_1 and P_2 , and for simplicity, we sometimes use x_1 and x_2 as sets of variable names.

$$\begin{aligned}
& \text{cstraces}^a(P_1 \parallel \text{cs} \parallel P_2) \\
& = \text{cstraces}^a(\text{begin state } [x_1 : T_1; x_2 : T_2] \bullet A_1 \parallel [x_1 \mid \text{cs} \mid x_2] A_2 \text{end}) \quad [\text{semantics of process parallelism}] \\
& = \text{cstraces}^a(w_1 \in T_1 \wedge w_2 \in T_2, x_1, x_2 := w_1, w_2, A_1 \parallel A_2) \\
& \quad \quad \quad [\text{definition of } \text{cstraces} \text{ for processes and abbreviation } \parallel \text{ for } [x_1 \mid \text{cs} \mid x_2]] \\
& = \{ \text{st}, \text{c}_2, \text{s}_2, \text{A}_3 \mid \alpha \text{st} \leq \mathbf{a} \wedge \quad [\text{definition of } \text{cstraces} \text{ for actions}] \\
& \quad \quad \quad (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1 \parallel A_2) \xrightarrow{\text{st}} (\text{c}_2 \mid \text{s}_2 \models \text{A}_3) \\
& \quad \bullet \mathcal{R}(\text{st}, \text{c}_2) \} \\
& = \{ \text{st}, \text{c}_2 \mid \alpha \text{st} \leq \mathbf{a} \wedge \quad [\text{Proposition 6}] \\
& \quad \exists \text{st}_1, \text{c}'_1, \text{s}'_1, \text{A}'_1, \text{st}_2, \text{c}'_2, \text{s}'_2, \text{A}'_2, \mathbf{b} \bullet \\
& \quad \quad (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{\text{st}_1} (\text{c}'_1 \mid \text{s}'_1 \models \text{A}'_1) \wedge \\
& \quad \quad (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_2) \xrightarrow{\text{st}_2} (\text{c}'_2 \mid \text{s}'_2 \models \text{A}'_2) \wedge \\
& \quad \quad \alpha \text{c}'_1 \cap \alpha \text{c}'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha \text{c}'_1 \wedge \alpha \text{st}_2 \subseteq \alpha \text{c}'_2 \wedge \alpha \text{st} \leq \mathbf{b} \wedge \\
& \quad \quad \mathcal{R}(\text{st}, \text{c}_2) \in \mathcal{N}^b (\mathcal{R} (\text{st}_1 \parallel \text{cs}, (\text{c}'_1 \wedge \text{c}'_2) \parallel \text{st}_2)) \\
& \quad \bullet \mathcal{R}(\text{st}, \text{c}_2) \} \\
& = \{ \text{st}, \text{c}_2 \mid \exists \text{st}_1, \text{c}'_1, \text{s}'_1, \text{A}'_1, \text{st}_2, \text{c}'_2, \text{s}'_2, \text{A}'_2 \bullet \quad [\text{Lemma 13 and } \alpha \text{st} \leq \mathbf{a} \text{ by definition of } \mathcal{N}^a \text{ and } \mathcal{R}] \\
& \quad \quad (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{\text{st}_1} (\text{c}'_1 \mid \text{s}'_1 \models \text{A}'_1) \wedge \\
& \quad \quad (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_2) \xrightarrow{\text{st}_2} (\text{c}'_2 \mid \text{s}'_2 \models \text{A}'_2) \wedge \\
& \quad \quad \alpha \text{c}'_1 \cap \alpha \text{c}'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha \text{c}'_1 \wedge \alpha \text{st}_2 \subseteq \alpha \text{c}'_2 \\
& \quad \quad \mathcal{R}(\text{st}, \text{c}_2) \in \mathcal{N}^a (\mathcal{R} (\text{st}_1 \parallel \text{cs}, (\text{c}'_1 \wedge \text{c}'_2) \parallel \text{st}_2)) \\
& \quad \bullet \mathcal{R}(\text{st}, \text{c}_2) \}
\end{aligned}$$

$$\begin{aligned}
&= \{ \text{st}, c_2 \mid \exists \text{st}_1, c'_1, s'_1, A'_1, \text{st}_2, c'_2, s'_2, A'_2 \bullet \quad \text{[Proposition 1]} \\
&\quad (\mathbf{w}_1 \in T_1 \wedge \mathbf{w}_2 \in T_2 \mid x_1 := \mathbf{w}_1 \models A_1) \xrightarrow{\text{st}_1} (c'_1 \mid s'_1 \models A'_1) \wedge \\
&\quad (\mathbf{w}_1 \in T_1 \wedge \mathbf{w}_2 \in T_2 \mid x_2 := \mathbf{w}_2 \models A_2) \xrightarrow{\text{st}_2} (c'_2 \mid s'_2 \models A'_2) \wedge \\
&\quad \alpha c'_1 \cap \alpha c'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha c'_1 \wedge \alpha \text{st}_2 \subseteq \alpha c'_2 \wedge \\
&\quad \mathcal{R}(\text{st}, c_2) \in \mathcal{N}^a (\mathcal{R} (\text{st}_1 \llbracket \text{cs}, (c'_1 \wedge c'_2) \rrbracket \text{st}_2)) \\
&\quad \bullet \mathcal{R}(\text{st}, c_2) \} \\
&= \{ \text{st}, c_2 \mid \exists \text{st}_1, c'_1, s'_1, A'_1, \text{st}_2, c'_2, s'_2, A'_2 \bullet \quad \text{[Proposition 2]} \\
&\quad (\mathbf{w}_1 \in T_1 \mid x_1 := \mathbf{w}_1 \models A_1) \xrightarrow{\text{st}_1} (c'_1 \mid s'_1 \models A'_1) \wedge \\
&\quad (\mathbf{w}_2 \in T_2 \mid x_2 := \mathbf{w}_2 \models A_2) \xrightarrow{\text{st}_2} (c'_2 \mid s'_2 \models A'_2) \wedge \\
&\quad \alpha c'_1 \cap \alpha c'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha c'_1 \wedge \alpha \text{st}_2 \subseteq \alpha c'_2 \wedge \\
&\quad \mathcal{R}(\text{st}, c_2) \in \mathcal{N}^a (\mathcal{R} (\text{st}_1 \llbracket \text{cs}, (c'_1 \wedge c'_2) \rrbracket \text{st}_2)) \\
&\quad \bullet \mathcal{R}(\text{st}, c_2) \} \\
&= \{ \text{st}, c_2 \mid \exists \text{st}_1, c'_1, \text{st}_2, c'_2, \mathbf{b}_1, \mathbf{b}_2 \bullet \quad \text{[definition of } cstraces] \\
&\quad \mathcal{R}(\text{st}_1, c'_1) \in cstraces^{\mathbf{b}_1}(\mathbf{w}_1 \in T_1, x_1 := \mathbf{w}_1, A_1) \wedge \\
&\quad \mathcal{R}(\text{st}_2, c'_2) \in cstraces^{\mathbf{b}_2}(\mathbf{w}_1 \in T_1, x_2 := \mathbf{w}_2, A_2) \wedge \\
&\quad \alpha c'_1 \cap \alpha c'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha c'_1 \wedge \alpha \text{st}_2 \subseteq \alpha c'_2 \wedge \\
&\quad \mathcal{R}(\text{st}, c_2) \in \mathcal{N}^a (\mathcal{R} (\text{st}_1 \llbracket \text{cs}, (c'_1 \wedge c'_2) \rrbracket \text{st}_2)) \\
&\quad \bullet \mathcal{R}(\text{st}, c_2) \} \\
&= \{ \text{st}, c_2 \mid \exists \text{st}_1, c'_1, \text{st}_2, c'_2, \mathbf{b}_1, \mathbf{b}_2 \bullet \quad \text{[Lemma 4, with } \mathbf{x} = \alpha(c'_1, c'_2) \setminus \alpha(\text{st}_1, \text{st}_2)] \\
&\quad \mathcal{R}(\text{st}_1, c'_1) \in cstraces^{\mathbf{b}_1}(\mathbf{w}_1 \in T_1, x_1 := \mathbf{w}_1, A_1) \wedge \\
&\quad \mathcal{R}(\text{st}_2, c'_2) \in cstraces^{\mathbf{b}_2}(\mathbf{w}_1 \in T_1, x_2 := \mathbf{w}_2, A_2) \wedge \\
&\quad \alpha c'_1 \cap \alpha c'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha c'_1 \wedge \alpha \text{st}_2 \subseteq \alpha c'_2 \wedge \\
&\quad \mathcal{R}(\text{st}, c_2) \in \mathcal{N}^a (\mathcal{R} (\text{st}_1 \llbracket \text{cs}, (\exists \mathbf{x} \bullet c'_1 \wedge c'_2) \rrbracket \text{st}_2)) \\
&\quad \bullet \mathcal{R}(\text{st}, c_2) \} \\
&= \{ \text{st}, c_2 \mid \exists \text{st}_1, c'_1, \text{st}_2, c'_2, \mathbf{b}_1, \mathbf{b}_2 \bullet \quad \text{[} \alpha c'_1 \cap \alpha c'_2 = \emptyset, \text{ with } \mathbf{x}_1 = \alpha c'_1 \setminus \alpha \text{st}_1 \text{ and } \mathbf{x}_2 = \alpha c'_2 \setminus \alpha \text{st}_2] \\
&\quad \mathcal{R}(\text{st}_1, c'_1) \in cstraces^{\mathbf{b}_1}(\mathbf{w}_1 \in T_1, x_1 := \mathbf{w}_1, A_1) \wedge \\
&\quad \mathcal{R}(\text{st}_2, c'_2) \in cstraces^{\mathbf{b}_2}(\mathbf{w}_1 \in T_1, x_2 := \mathbf{w}_2, A_2) \wedge \\
&\quad \alpha c'_1 \cap \alpha c'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha c'_1 \wedge \alpha \text{st}_2 \subseteq \alpha c'_2 \wedge \\
&\quad \mathcal{R}(\text{st}, c_2) \in \mathcal{N}^a (\mathcal{R} (\text{st}_1 \llbracket \text{cs}, ((\exists \mathbf{x}_1 \bullet c'_1) \wedge (\exists \mathbf{x}_2 \bullet c'_2)) \rrbracket \text{st}_2)) \\
&\quad \bullet \mathcal{R}(\text{st}, c_2) \} \\
&= \{ \text{st}, c_2 \mid \exists \text{st}_1, c'_1, \text{st}_2, c'_2, \mathbf{b}_1, \mathbf{b}_2 \bullet \quad \text{[definition of synchronisation operator]} \\
&\quad \mathcal{R}(\text{st}_1, c'_1) \in cstraces^{\mathbf{b}_1}(\mathbf{w}_1 \in T_1, x_1 := \mathbf{w}_1, A_1) \wedge \\
&\quad \mathcal{R}(\text{st}_2, c'_2) \in cstraces^{\mathbf{b}_2}(\mathbf{w}_1 \in T_1, x_2 := \mathbf{w}_2, A_2) \wedge \\
&\quad \alpha c'_1 \cap \alpha c'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha c'_1 \wedge \alpha \text{st}_2 \subseteq \alpha c'_2 \wedge \\
&\quad \mathcal{R}(\text{st}, c_2) \in \mathcal{N}^a (\text{st}_1, \exists \mathbf{x}_1 \bullet c'_1) \llbracket \text{cs} \rrbracket (\text{st}_2, \exists \mathbf{x}_2 \bullet c'_2) \\
&\quad \bullet \mathcal{R}(\text{st}, c_2) \} \\
&= \{ \text{st}, c_2 \mid \exists \text{st}_1, c'_1, \text{st}_2, c'_2, \mathbf{b}_1, \mathbf{b}_2 \bullet \quad \text{[definition of } \mathcal{R}] \\
&\quad \mathcal{R}(\text{st}_1, c'_1) \in cstraces^{\mathbf{b}_1}(\mathbf{w}_1 \in T_1, x_1 := \mathbf{w}_1, A_1) \wedge \\
&\quad \mathcal{R}(\text{st}_2, c'_2) \in cstraces^{\mathbf{b}_2}(\mathbf{w}_1 \in T_1, x_2 := \mathbf{w}_2, A_2) \wedge \\
&\quad \alpha c'_1 \cap \alpha c'_2 = \emptyset \wedge \alpha \text{st}_1 \subseteq \alpha c'_1 \wedge \alpha \text{st}_2 \subseteq \alpha c'_2 \wedge \\
&\quad \mathcal{R}(\text{st}, c_2) \in \mathcal{N}^a (\mathcal{R}(\text{st}_1, c'_1) \llbracket \text{cs} \rrbracket \mathcal{R}(\text{st}_2, c'_2)) \\
&\quad \bullet \mathcal{R}(\text{st}, c_2) \} \\
&= \{ \text{st}, c_2 \mid \exists \mathbf{b}_1, \mathbf{b}_2 \bullet \quad \text{[predicate calculus and property of } cstraces] \\
&\quad \exists \text{cst}_1 : cstraces^{\mathbf{b}_1}(\mathbf{w}_1 \in T_1, x_1 := \mathbf{w}_1, A_1); \text{cst}_2 : cstraces^{\mathbf{b}_2}(\mathbf{w}_1 \in T_1, x_2 := \mathbf{w}_2, A_2) \bullet \\
&\quad \text{disjoint}(\mathbf{b}_1, \mathbf{b}_2) \wedge \mathcal{R}(\text{st}, c_2) \in \mathcal{N}^a (\text{cst}_1 \llbracket \text{cs} \rrbracket \text{cst}_2) \\
&\quad \bullet \mathcal{R}(\text{st}, c_2) \}
\end{aligned}$$

$$\begin{aligned}
&= \bigcup \{ \mathbf{b}_1, \mathbf{b}_2, \mathbf{cst}_1 : \mathit{cstraces}^{\mathbf{b}_1}(\mathbf{w}_1 \in \mathbf{T}_1, \mathbf{x}_1 := \mathbf{w}_1, \mathbf{A}_1); \mathbf{cst}_2 : \mathit{cstraces}^{\mathbf{b}_2}(\mathbf{w}_1 \in \mathbf{T}_1, \mathbf{x}_2 := \mathbf{w}_2, \mathbf{A}_2) \mid \\
&\quad \text{disjoint}(\mathbf{b}_1, \mathbf{b}_2) \\
&\quad \bullet \mathcal{N}^a(\llbracket \mathbf{cst}_1 \parallel \mathbf{cs} \rrbracket \mathbf{cst}_2) \} \\
&\hspace{20em} [\text{property of sets and relational image}] \\
&= \bigcup \{ \mathbf{cst}_1 : \mathit{cstraces}^{\mathbf{b}_1}(\mathbf{w}_1 \in \mathbf{T}_1, \mathbf{x}_1 := \mathbf{w}_1, \mathbf{A}_1); \mathbf{cst}_2 : \mathit{cstraces}^{\mathbf{b}_2}(\mathbf{w}_1 \in \mathbf{T}_1, \mathbf{x}_2 := \mathbf{w}_2, \mathbf{A}_2) \\
&\quad \bullet \mathcal{N}^a(\llbracket \mathbf{cst}_1 \parallel \mathbf{cs} \rrbracket \mathbf{cst}_2) \} \\
&\hspace{20em} [\text{Lemma 12 and assumption: disjoint}(\mathbf{b}_1, \mathbf{b}_2)] \\
&\hspace{20em} \square
\end{aligned}$$

Appendix C.2. Properties of satisfiability

Satisfiability, as captured by $\mathit{satisfiable}_{\mathbf{cs}}((\mathbf{st}_1, \mathbf{c}_1), (\mathbf{st}_2, \mathbf{c}_2))$, does not ensure that there is an actual synchronisation arising from the traces \mathbf{st}_1 and \mathbf{st}_2 , since their projections $(\mathbf{st}_1 \upharpoonright_{\mathbf{E} \mathbf{cs}})$ and $(\mathbf{st}_2 \upharpoonright_{\mathbf{E} \mathbf{cs}})$ might be empty. It does ensure, however, that the set of traces defined by the synchronisation operator is not empty. This is established by the Lemma 10 proved below.

Lemma 10.

$$(\mathbf{st}_1 \parallel \mathbf{cs}, (\mathbf{c}_1 \wedge \mathbf{c}_2) \parallel \mathbf{st}_2) \neq \emptyset \Leftrightarrow \mathit{satisfiable}_{\mathbf{cs}}((\mathbf{st}_1, \mathbf{c}_1), (\mathbf{st}_2, \mathbf{c}_2))$$

Proof.

$$\begin{aligned}
&(\mathbf{st}_1 \parallel \mathbf{cs}, (\mathbf{c}_1 \wedge \mathbf{c}_2) \parallel \mathbf{st}_2) \neq \emptyset \\
&\Leftrightarrow \exists \mathbf{a}_1, \mathbf{a}_2 \bullet (\mathbf{st}_1 \upharpoonright_{\mathbf{E} \mathbf{cs}}) =_{\mathbf{S}} (\mathbf{st}_2 \upharpoonright_{\mathbf{E} \mathbf{cs}}) \wedge \mathbf{c}_1 \wedge \mathbf{c}_2 \hspace{10em} [\text{Lemma 5, with } \alpha \mathbf{st}_1 \leq \mathbf{a}_1 \text{ and } \alpha \mathbf{st}_2 \leq \mathbf{a}_2] \\
&\Leftrightarrow \mathit{satisfiable}_{\mathbf{cs}}((\mathbf{st}_1, \mathbf{c}_1), (\mathbf{st}_2, \mathbf{c}_2)) \hspace{10em} [\text{definition of } \mathit{satisfiable}] \\
&\hspace{20em} \square
\end{aligned}$$

This relationship between the notion of satisfiability and the existence of synchronisation traces reassures us of the adequacy of the definition of satisfiability.

The following lemma identifies a stronger condition that holds when the set of synchronisation traces not only is not empty, but also contains at least one trace with at least one synchronisation.

Lemma 11.

$$\begin{aligned}
&\exists \mathbf{st}, \mathbf{c} \bullet (\mathbf{st}, \mathbf{c}) \in (\mathbf{st}_1 \parallel \mathbf{cs}, (\mathbf{c}_1 \wedge \mathbf{c}_2) \parallel \mathbf{st}_2) \wedge \mathbf{st} \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \\
&\Rightarrow \\
&\mathbf{st}_1 \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \wedge \mathbf{st}_2 \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \wedge \mathit{satisfiable}_{\mathbf{cs}}((\mathbf{st}_1, \mathbf{c}_1), (\mathbf{st}_2, \mathbf{c}_2))
\end{aligned}$$

Proof. By induction.

Case 1.

$$\begin{aligned}
&\exists \mathbf{st}, \mathbf{c} \bullet (\mathbf{st}, \mathbf{c}) \in (\langle \rangle \parallel \mathbf{cs}, (\mathbf{c}_1 \wedge \mathbf{c}_2) \parallel \langle \rangle) \wedge \mathbf{st} \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \\
&\Leftrightarrow \text{false} \hspace{15em} [\text{definitions of } (\langle \rangle \parallel \mathbf{cs}, \mathbf{c}_2 \parallel \langle \rangle) \text{ and } \upharpoonright_{\mathbf{C}}] \\
&\Leftrightarrow \langle \rangle \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \wedge \langle \rangle \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \wedge \mathit{satisfiable}_{\mathbf{cs}}((\langle \rangle, \mathbf{c}_1), (\langle \rangle, \mathbf{c}_2)) \hspace{10em} [\langle \rangle \upharpoonright_{\mathbf{C} \mathbf{cs}} = \langle \rangle]
\end{aligned}$$

Case 2. : $\mathbf{d} \notin \mathbf{cs}$

$$\begin{aligned}
&\exists \mathbf{st}, \mathbf{c} \bullet (\mathbf{st}, \mathbf{c}) \in (\langle \rangle \parallel \mathbf{cs}, (\mathbf{c}_1 \wedge \mathbf{c}_2) \parallel \langle \mathbf{d}, \beta \rangle) \wedge \mathbf{st} \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \\
&\Leftrightarrow \text{false} \hspace{15em} [\text{definition of } (\langle \rangle \parallel \mathbf{cs}, \mathbf{c}_2 \parallel \langle \mathbf{d}, \beta \rangle) \text{ and } \langle \mathbf{d}, \beta \rangle \upharpoonright_{\mathbf{C} \mathbf{cs}} = \langle \rangle, \text{ if } \mathbf{d} \notin \mathbf{cs}] \\
&\Leftrightarrow \langle \rangle \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \wedge \langle \mathbf{d}, \beta \rangle \upharpoonright_{\mathbf{C} \mathbf{cs}} \neq \langle \rangle \wedge \mathit{satisfiable}_{\mathbf{cs}}((\langle \rangle, \mathbf{c}_1), (\langle \mathbf{d}, \beta \rangle, \mathbf{c}_2)) \quad [\langle \rangle \upharpoonright_{\mathbf{C} \mathbf{cs}} = \langle \rangle \text{ (and } \langle \mathbf{d}, \beta \rangle \upharpoonright_{\mathbf{C} \mathbf{cs}} = \langle \rangle)]
\end{aligned}$$

Case 3: $d_1 \in cs, d_2 \notin cs$

$$\begin{aligned}
& (\exists st, c \bullet (st, c) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, (c_1 \wedge c_2) \rrbracket ((d_2.\beta) \wedge st_2)) \wedge st \upharpoonright_C cs \neq \langle \rangle) \\
& \Leftrightarrow (\exists st, c \bullet (st, c) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, (c_1 \wedge c_2) \rrbracket st_2) \wedge ((d_2.\beta) \wedge st) \upharpoonright_C cs \neq \langle \rangle) \\
& \hspace{15em} [\text{definition of } ((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d_2.\beta) \wedge st_2)] \\
& \Leftrightarrow \exists st, c \bullet (st, c) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, (c_1 \wedge c_2) \rrbracket st_2) \wedge st \upharpoonright_C cs \neq \langle \rangle \hspace{10em} [\text{definition of } \upharpoonright_C] \\
& \Leftrightarrow ((d_1.\alpha) \wedge st_1) \upharpoonright_C cs \neq \langle \rangle \wedge st_2 \upharpoonright_C cs \neq \langle \rangle \wedge \text{satisfiable}_{cs}(((d_1.\alpha) \wedge st_1, c_1), (st_2, c_2)) \\
& \hspace{15em} [\text{induction hypothesis}] \\
& \Leftrightarrow ((d_1.\alpha) \wedge st_1) \upharpoonright_C cs \neq \langle \rangle \wedge ((d_2.\alpha) \wedge st_2) \upharpoonright_C cs \neq \langle \rangle \wedge \text{satisfiable}_{cs}(((d_1.\alpha) \wedge st_1, c_1), (st_2, c_2)) \\
& \hspace{15em} [\text{definition of } \upharpoonright_C] \\
& \Leftrightarrow ((d_1.\alpha) \wedge st_1) \upharpoonright_C cs \neq \langle \rangle \wedge ((d_2.\alpha) \wedge st_2) \upharpoonright_C cs \neq \langle \rangle \wedge \text{satisfiable}_{cs}(((d_1.\alpha) \wedge st_1, c_1), ((d_2.\alpha) \wedge st_2, c_2)) \\
& \hspace{15em} [\text{definitions of } \text{satisfiable} \text{ and } \upharpoonright_C]
\end{aligned}$$

Case 4: $d \in cs$

$$\begin{aligned}
& \exists st, c \bullet (st, c) \in (((d.\alpha) \wedge st_1) \llbracket cs, (c_1 \wedge c_2) \rrbracket ((d.\beta) \wedge st_2)) \wedge st \upharpoonright_C cs \neq \langle \rangle \\
& \Leftrightarrow (\exists st, c_1 \bullet (st, c) \in (st_1 \llbracket cs, (c_1 \wedge c_2 \wedge \alpha = \beta) \rrbracket st_2) \wedge ((d.\beta) \wedge st) \upharpoonright_C cs \neq \langle \rangle) \\
& \hspace{15em} [\text{definition of } ((d.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d.\beta) \wedge st_2)] \\
& \Leftrightarrow \exists st, c_1 \bullet (st, c) \in (st_1 \llbracket cs, (c_1 \wedge c_2 \wedge \alpha = \beta) \rrbracket st_2) \hspace{5em} [((d.\beta) \wedge st) \upharpoonright_C cs \neq \langle \rangle \text{ by definition of } \upharpoonright_C] \\
& \Leftrightarrow \exists a_1, a_2 \bullet (st_1 \upharpoonright_E cs) =_S (st_2 \upharpoonright_E cs) \wedge c_1 \wedge c_2 \wedge \alpha = \beta \hspace{5em} [\text{Lemma 5, with } \alpha st_1 \leq a_1 \text{ and } \alpha st_2 \leq a_2] \\
& \Leftrightarrow \exists a_1, a_2 \bullet (((d.\alpha) \wedge st_1) \upharpoonright_E cs) =_S (((d.\beta) \wedge st_2) \upharpoonright_E cs) \wedge c_1 \wedge c_2 \hspace{5em} [\text{definitions of } \upharpoonright_E \text{ and } =_S] \\
& \Leftrightarrow \text{satisfiable}_{cs}(((d.\alpha) \wedge st_1, c_1), ((d.\beta) \wedge st_2, c_2)) \hspace{10em} [\text{definition of } \text{satisfiable}] \\
& \Leftrightarrow ((d.\alpha) \wedge st_1) \upharpoonright_C cs \neq \langle \rangle \wedge ((d.\beta) \wedge st_2) \upharpoonright_C cs \neq \langle \rangle \wedge \text{satisfiable}_{cs}(((d.\alpha) \wedge st_1, c_1), ((d.\beta) \wedge st_2, c_2)) \\
& \hspace{15em} [\text{definition of } \upharpoonright_C]
\end{aligned}$$

Case 5: $d_1 \neq d_2, d_1 \in cs, d_2 \in cs$

$$\begin{aligned}
& (\exists st, c \bullet (st, c) \in (((d_1.\alpha) \wedge st_1) \llbracket cs, (c_1 \wedge c_2) \rrbracket ((d_2.\beta) \wedge st_2)) \wedge st \upharpoonright_C cs \neq \langle \rangle) \\
& \Leftrightarrow \text{false} \hspace{15em} [((d_1.\alpha) \wedge st_1) \llbracket cs, c_2 \rrbracket ((d_2.\beta) \wedge st_2) = \emptyset] \\
& \Leftrightarrow ((d_1.\alpha) \wedge st_1) \upharpoonright_C cs \neq \langle \rangle \wedge ((d_2.\beta) \wedge st_2) \upharpoonright_C cs \neq \langle \rangle \wedge \text{satisfiable}_{cs}(((d_1.\alpha) \wedge st_1, c_1), ((d_2.\beta) \wedge st_2, c_2)) \\
& \hspace{15em} [\text{definition of } \text{satisfiable} \text{ and } ((d_1.\alpha) \wedge st_1 \upharpoonright_E cs) =_S ((d_2.\beta) \wedge st_2 \upharpoonright_E cs) \text{ is false}]
\end{aligned}$$

Case 6.: $d_1 \notin cs, d_2 \notin cs$

$$\begin{aligned}
& (\exists st, c \bullet (st, c) \in ((\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, (c_1 \wedge c_2) \rrbracket (\langle d_2.\beta \rangle \wedge st_2)) \wedge st \upharpoonright_C cs \neq \langle \rangle) \\
& \Leftrightarrow (\exists st, c \bullet (st, c_1) \in ((\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, (c_1 \wedge c_2) \rrbracket st_2) \wedge (\langle d_2.\beta \rangle \wedge st) \upharpoonright_C cs \neq \langle \rangle) \\
& \quad \vee \\
& \quad (\exists st, c \bullet (st, c) \in (st_1 \llbracket cs, (c_1 \wedge c_2) \rrbracket (\langle d_2.\beta \rangle \wedge st_2)) \wedge (\langle d_1.\alpha \rangle \wedge st) \upharpoonright_C cs \neq \langle \rangle) \\
& \hspace{15em} [\text{definition of } ((\langle d_1.\alpha \rangle \wedge st_1) \llbracket cs, c_2 \rrbracket (\langle d_2.\beta \rangle \wedge st_2))] \\
& \Leftrightarrow ((\langle d_1.\alpha \rangle \wedge st_1) \upharpoonright_C cs \neq \langle \rangle \wedge st_2 \upharpoonright_C cs \neq \langle \rangle \wedge \text{satisfiable}_{cs}((\langle d_1.\alpha \rangle \wedge st_1, c_1), (st_2, c_2))) \vee \\
& \quad (st_1 \upharpoonright_C cs \neq \langle \rangle \wedge (\langle d_2.\beta \rangle \wedge st_2) \upharpoonright_C cs \neq \langle \rangle \wedge \text{satisfiable}_{cs}((st_1, c_1), (\langle d_2.\beta \rangle \wedge st_2, c_2))) \\
& \hspace{15em} [\text{induction hypothesis}] \\
& \Leftrightarrow ((\langle d_1.\alpha \rangle \wedge st_1) \upharpoonright_C cs \neq \langle \rangle \wedge (\langle d_2.\beta \rangle \wedge st_2) \upharpoonright_C cs \neq \langle \rangle \wedge \\
& \quad \text{satisfiable}_{cs}((\langle d_1.\alpha \rangle \wedge st_1, c_1), (\langle d_2.\beta \rangle \wedge st_2, c_2)) \\
& \hspace{10em} [(\langle d_2.\beta \rangle \wedge st_2) \upharpoonright_C cs = st_2 \upharpoonright_C cs \text{ and } (\langle d_2.\beta \rangle \wedge st_2) \upharpoonright_E cs = st_2 \upharpoonright_E cs, \text{ and}] \\
& \hspace{10em} [(\langle d_1.\alpha \rangle \wedge st_1) \upharpoonright_C cs = st_1 \upharpoonright_C cs \text{ and } (\langle d_1.\alpha \rangle \wedge st_1) \upharpoonright_E cs = st_1 \upharpoonright_E cs] \\
& \hspace{15em} \square
\end{aligned}$$

The reverse implication is a direct consequence of Lemmas 10 and 6.

Appendix C.3. Properties of normalisation

When applied to equivalent constrained symbolic traces, \mathcal{N} makes them syntactically equal.

Lemma 12.

$$cst_1 \equiv cst_2 \Rightarrow \mathcal{N}^a(cst_1) = \mathcal{N}^a(cst_2)$$

Proof. Direct from the definition of \mathcal{N} and equivalence of constrained symbolic traces: $(st_1, c_1) \equiv (st_2, c_2)$ if and only if $st_1 \upharpoonright_C = st_2 \upharpoonright_C$ and $c_1[\alpha st_2 / \alpha st_1] = c_2$. The projection operation $st \upharpoonright_C$ keeps just the channels of the trace st . \square

The following lemma gives a property of alphabets and \mathcal{N} .

Lemma 13. For every a_1 and a_2 such that $\alpha st \leq a_1$ and $\alpha st \leq a_2$,

$$(st, c) = \mathcal{N}^{a_1}(cst) \Leftrightarrow (st, c) = \mathcal{N}^{a_2}(cst)$$

Proof.

$$\begin{aligned}
& (st, c) = \mathcal{N}^{a_1}(st_1, c_1) \\
& \Leftrightarrow (st = st_1[a_1 / \alpha st]) \wedge (c = c_1[a_1 / \alpha st]) \hspace{10em} [\text{definition of } \mathcal{N}^{a_1}] \\
& \Leftrightarrow (st = st_1[a_2 / \alpha st]) \wedge (c = c_1[a_2 / \alpha st]) \hspace{10em} [\alpha st \leq a_1 \text{ and } \alpha st \leq a_2 \text{ imply } (a_1 \upharpoonright \# st) \leq a_2] \\
& \Leftrightarrow (st, c) = \mathcal{N}^{a_2}(st_1, c_1) \hspace{10em} [\text{definition of } \mathcal{N}^{a_2}]
\end{aligned}$$

We use $(a \upharpoonright \# n)$ to refer to the first n elements of the alphabet a , and $\# st$ for the size of a trace st . \square

Appendix D. Proofs of theorems

i

Theorem 2.

$$Scstraces_{cs}^{a_1, (a_1, a_2)}(P_1, P_2) = \{ \text{cst} : cstraces^a(P_1 \llbracket cs \rrbracket P_2) \mid \text{cst} \upharpoonright_C cs \neq \langle \rangle \}$$

Proof.

$$\begin{aligned} & Scstraces_{cs}^{a_1, (a_1, a_2)}(P_1, P_2) \\ &= \bigcup \{ p : ScstraceP_{cs}^{(a_1, a_2)}(P_1, P_2) \bullet \mathcal{N}^a(\langle p.1 \llbracket cs \rrbracket p.2 \rangle) \} && \text{[definition of } Scstraces \text{]} \\ &= \bigcup \{ \text{cst}_1 : cstraces^{a_1}(P_1); \text{cst}_2 : cstraces^{a_2}(P_2) \mid && \text{[definition of } ScstraceP \text{]} \\ &\quad (\text{cst}_1 \upharpoonright_C cs) \neq \langle \rangle \wedge (\text{cst}_2 \upharpoonright_C cs) \neq \langle \rangle \wedge \text{satisfiable}_{cs}(\text{cst}_1, \text{cst}_2) \\ &\quad \bullet \mathcal{N}^a(\langle \text{cst}_1 \llbracket cs \rrbracket \text{cst}_2 \rangle) \} \\ &= \{ \text{cst}_1 : cstraces^{a_1}(P_1); \text{cst}_2 : cstraces^{a_2}(P_2); \text{cst} : \mathcal{N}^a(\langle \text{cst}_1 \llbracket cs \rrbracket \text{cst}_2 \rangle) \mid && \text{[property of sets]} \\ &\quad (\text{cst}_1 \upharpoonright_C cs) \neq \langle \rangle \wedge (\text{cst}_2 \upharpoonright_C cs) \neq \langle \rangle \wedge \text{satisfiable}_{cs}(\text{cst}_1, \text{cst}_2) \\ &\quad \bullet \text{cst} \} \\ &= \{ \text{cst}_1 : cstraces^{a_1}(P_1); \text{cst}_2 : cstraces^{a_2}(P_2); \text{cst} : \mathcal{N}^a(\langle \text{cst}_1 \llbracket cs \rrbracket \text{cst}_2 \rangle) \mid && \text{[Lemmas 10 and 7]} \\ &\quad (\text{cst}_1 \upharpoonright_C cs) \neq \langle \rangle \wedge (\text{cst}_2 \upharpoonright_C cs) \neq \langle \rangle \wedge \text{satisfiable}_{cs}(\text{cst}_1, \text{cst}_2) \wedge (\text{cst} \upharpoonright_C cs) \neq \langle \rangle \\ &\quad \bullet \text{cst} \} \\ &= \{ \text{cst}_1 : cstraces^{a_1}(P_1); \text{cst}_2 : cstraces^{a_2}(P_2); \text{cst} : \mathcal{N}^a(\langle \text{cst}_1 \llbracket cs \rrbracket \text{cst}_2 \rangle) \mid (\text{cst} \upharpoonright_C cs) \neq \langle \rangle && \text{[Lemma 8]} \\ &\quad \bullet \text{cst} \} \\ &= \{ \text{cst} : \bigcup \{ \text{cst}_1 : cstraces^{a_1}(P_1); \text{cst}_2 : cstraces^{a_2}(P_2) \bullet \mathcal{N}^a(\langle \text{cst}_1 \llbracket cs \rrbracket \text{cst}_2 \rangle) \} \mid && \text{[property of sets]} \\ &\quad (\text{cst} \upharpoonright_C cs) \neq \langle \rangle \} \\ &= \{ \text{cst} : cstraces^a(P_1 \llbracket cs \rrbracket P_2) \mid \text{cst} \upharpoonright_C cs \neq \langle \rangle \} && \text{[Lemma 9]} \end{aligned}$$

□

Theorem 3.

$$\begin{aligned} & Scstraces_{cs}^a(Scstraces_{cs}^{a_{12}}(cstraces^{a_1}(P_1), cstraces^{a_2}(P_2)), cstraces^{a_3}(P_3)) \\ &= \\ & Scstraces_{cs}^a(cstraces^{a_{12}}(P_1 \llbracket cs \rrbracket P_2), cstraces^{a_3}(P_3)) \end{aligned}$$

Proof.

$$\begin{aligned} & Scstraces_{cs}^a(Scstraces_{cs}^{a_{12}}(cstraces^{a_1}(P_1), cstraces^{a_2}(P_2)), cstraces^{a_3}(P_3)) \\ &= Scstraces_{cs}^a(\bigcup \{ p : ScstraceP_{cs}(cstraces^{a_1}(P_1), cstraces^{a_2}(P_2)) \bullet \mathcal{N}^{a_{12}}(\langle p.1 \llbracket cs \rrbracket p.2 \rangle) \}, cstraces^{a_3}(P_3)) && \text{[Definition 25]} \\ &= Scstraces_{cs}^a(\bigcup \{ \text{cst}_1 : cstraces^{a_1}(P_1); \text{cst}_2 : cstraces^{a_2}(P_2) \mid && \text{[Definition 24]} \\ &\quad (\text{cst}_1 \upharpoonright_C cs) \neq \langle \rangle \wedge (\text{cst}_2 \upharpoonright_C cs) \neq \langle \rangle \wedge \text{satisfiable}_{cs}(\text{cst}_1, \text{cst}_2) \\ &\quad \bullet \mathcal{N}^{a_{12}}(\langle \text{cst}_1 \llbracket cs \rrbracket \text{cst}_2 \rangle) \}, \\ &\quad cstraces^{a_3}(P_3)) \\ &= Scstraces_{cs}^a(\bigcup \{ p : ScstraceP_{cs}^{(a_1, a_2)}(P_1, P_2) \bullet \mathcal{N}^{a_{12}}(\langle p.1 \llbracket cs \rrbracket p.2 \rangle) \}, cstraces^{a_3}(P_3)) && \text{[Definition 12]} \\ &= Scstraces_{cs}^a(Scstraces_{cs}^{a_{12}, (a_1, a_2)}(P_1, P_2), cstraces^{a_3}(P_3)) && \text{[Definition 15]} \\ &= Scstraces_{cs}^a(\{ \text{cst} : cstraces^{a_{12}}(P_1 \llbracket cs \rrbracket P_2) \mid \text{cst} \upharpoonright_C cs \neq \langle \rangle \}, cstraces^{a_3}(P_3)) && \text{[Theorem 2]} \\ &= \bigcup \{ p : ScstraceP_{cs}(\{ \text{cst} : cstraces^{a_{12}}(P_1 \llbracket cs \rrbracket P_2) \mid \text{cst} \upharpoonright_C cs \neq \langle \rangle \}, cstraces^{a_3}(P_3)) \bullet && \\ &\quad \mathcal{N}^a(\langle p.1 \llbracket cs \rrbracket p.2 \rangle) \} \end{aligned}$$

[definition of $Scstraces$]

$$\begin{aligned}
&= \bigcup \{ p : \{ \text{cst}_1 : \{ \text{cst} : cstraces^{a_{12}}(P_1 \llbracket \text{cs} \rrbracket P_2) \mid \text{cst} \downarrow_C \text{cs} \neq \langle \rangle \}; \text{cst}_2 : cstraces^{a_3}(P_3) \mid \\
&\quad (\text{cst}_1 \downarrow_C \text{cs}) \neq \langle \rangle \wedge (\text{cst}_2 \downarrow_C \text{cs}) \neq \langle \rangle \wedge \text{satisfiable}_{cs}(\text{cst}_1, \text{cst}_2) \} \bullet \\
&\quad \mathcal{N}^a(p.1 \llbracket \text{cs} \rrbracket p.2) \} \quad \text{[definition of } ScstraceP\text{]} \\
&= \bigcup \{ \text{cst}_1 : \{ \text{cst} : cstraces^{a_{12}}(P_1 \llbracket \text{cs} \rrbracket P_2) \mid \text{cst} \downarrow_C \text{cs} \neq \langle \rangle \}; \text{cst}_2 : cstraces^{a_3}(P_3) \mid \quad \text{[property of sets]} \\
&\quad (\text{cst}_1 \downarrow_C \text{cs}) \neq \langle \rangle \wedge (\text{cst}_2 \downarrow_C \text{cs}) \neq \langle \rangle \wedge \text{satisfiable}_{cs}(\text{cst}_1, \text{cst}_2) \bullet \\
&\quad \mathcal{N}^a(\text{cst}_1 \llbracket \text{cs} \rrbracket \text{cst}_2) \} \\
&= \bigcup \{ \text{cst}_1 : cstraces^{a_{12}}(P_1 \llbracket \text{cs} \rrbracket P_2); \text{cst}_2 : cstraces^{a_3}(P_3) \mid \quad \text{[property of sets]} \\
&\quad (\text{cst}_1 \downarrow_C \text{cs}) \neq \langle \rangle \wedge (\text{cst}_2 \downarrow_C \text{cs}) \neq \langle \rangle \wedge \text{satisfiable}_{cs}(\text{cst}_1, \text{cst}_2) \bullet \\
&\quad \mathcal{N}^a(\text{cst}_1 \llbracket \text{cs} \rrbracket \text{cst}_2) \} \\
&= \bigcup \{ p : ScstraceP_{cs}^{(a_{12}, a_3)}(P_1 \llbracket \text{cs} \rrbracket P_2, P_3) \mid \mathcal{N}^a(p.1 \llbracket \text{cs} \rrbracket p.2) \} \quad \text{[definition of } ScstraceP\text{]} \\
&= Scstraces_{cs}^{(a_{12}, a_3)}(P_1 \llbracket \text{cs} \rrbracket P_2, P_3) \quad \text{[definition of } Scstraces\text{]} \\
&= Scstraces_{cs}^a(cstraces^{a_{12}}(P_1 \llbracket \text{cs} \rrbracket P_2), cstraces^{a_3}(P_3)) \quad \text{[definition of } Scstraces\text{]}
\end{aligned}$$

□

Theorem 4.

$$selectS^a(P) \subseteq cstraces^a(P)$$

Proof. By induction.

Case: basic process. Trivial, because $selectS(P) = \emptyset$.

Case: parallelism. Direct consequence of Theorem 2.

Case: choices. We can assume, without loss of generality, that

$$P_1 = \text{begin state } [x_1 : T_1] \bullet A_1 \text{ end and } P_2 = \text{begin state } [x_2 : T_2] \bullet A_2 \text{ end}$$

with $x_1 \cap x_2 = \emptyset$, where x_1 and x_2 are the lists of state components of P_1 and P_2 , and for simplicity, we sometimes use x_1 and x_2 as sets of variable names.

$$\begin{aligned}
&selectS^a(P_1 \square P_2) \\
&= selectS^a(P_1) \cup selectS^a(P_2) \quad \text{[definition of } selectS\text{]} \\
&\subseteq cstraces^a(P_1) \cup cstraces^a(P_2) \quad \text{[induction hypothesis]} \\
&= cstraces^a(w_1 \in T_1, x_1 := w_1, A_1) \cup cstraces^a(w_2 \in T_2, x_2 := w_2, A_2) \quad \text{[definition of } cstraces\text{]} \\
&= cstraces^a(w_1 \in T_1 \wedge w_2 \in T_2, (x_1, x_2 := w_1, w_2), A_1) \cup cstraces^a(w_1 \in T_1 \wedge w_2 \in T_2, (x_1, x_2 := w_1, w_2), A_2) \\
&\quad \text{[definition of } cstraces\text{, and Propositions 1 and 2]} \\
&= \{ \text{st}, c_3, s_3, A_3 \mid \alpha \text{st} \leq a \wedge (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{\text{st}} (c_3 \mid s_3 \models A_3) \bullet \mathcal{R}(\text{st}, c_2) \} \\
&\quad \cup \\
&\quad \{ \text{st}, c_3, s_3, A_3 \mid \alpha \text{st} \leq a \wedge (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_2) \xrightarrow{\text{st}} (c_3 \mid s_3 \models A_3) \bullet \mathcal{R}(\text{st}, c_3) \}
\end{aligned}$$

[definition of *cstraces*]

$$= \{ \text{st}, c_3, s_3, A_3 \mid \begin{array}{l} \alpha \text{st} \leq a \wedge (\bar{w}_1 \in T_1 \wedge \bar{w}_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{\text{st}} (c_3 \mid s_3 \models A_3) \vee \\ \alpha \text{st} \leq a \wedge (\bar{w}_1 \in T_1 \wedge \bar{w}_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_2) \xrightarrow{\text{st}} (c_3 \mid s_3 \models A_3) \\ \bullet \mathcal{R}(\text{st}, c_3) \end{array} \} \quad \text{[property of sets]}$$

$$= \{ \text{st}, c_3, s_3, A_3 \mid \alpha \text{st} \leq a \wedge (\bar{w}_1 \in T_1 \wedge \bar{w}_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1 \sqcap A_2) \xrightarrow{\text{st}} (c_3 \mid s_3 \models A_3) \\ \bullet \mathcal{R}(\text{st}, c_3) \}$$

[Proposition 4]

$$= \text{cstraces}^a(\bar{w}_1 \in T_1 \wedge \bar{w}_2 \in T_2, (x_1, x_2 := w_1, w_2), A_1 \sqcap A_2) \quad \text{[definition of } cstraces]$$

$$= \text{cstraces}^a(P_1 \sqcap P_2) \quad \text{[definition of process external choice and } cstraces]$$

The proof is similar for internal choice.

Case: sequence. We can assume, without loss of generality, that

$$P_1 = \text{begin state } [x_1 : T_1] \bullet A_1 \text{ end and } P_2 = \text{begin state } [x_2 : T_2] \bullet A_2 \text{ end}$$

with $x_1 \cap x_2 = \emptyset$, where x_1 and x_2 are the lists of state components of P_1 and P_2 , and for simplicity, we sometimes use x_1 and x_2 as sets of variable names.

$$\text{selectS}^a(P_1 ; P_2)$$

$$= \text{selectS}^a(P_1) \cup \mathcal{N}^a(\text{tcstraces}^{a_1}(P_1) \text{ ccat } \text{selectS}^{a_2}(P_2)) \quad \text{[definition of } selectS]$$

$$\subseteq \text{cstraces}^a(P_1) \cup \mathcal{N}^a(\text{tcstraces}^{a_1}(P_1) \text{ ccat } \text{cstraces}^{a_2}(P_2)) \quad \text{[induction hypothesis]}$$

$$= \text{cstraces}^a(\bar{w}_1 \in T_1, x_1 := w_1, A_1) \cup \mathcal{N}^a(\text{tcstraces}^{a_1}(\bar{w}_1 \in T_1, x_1 := w_1, A_1) \text{ ccat } \text{cstraces}^{a_2}(\bar{w}_2 \in T_2, x_2 := w_2, A_2)) \quad \text{[definitions of } cstraces \text{ and } tcstraces]$$

$$= \{ \text{st}, c_3, s_3, A_3 \mid \alpha \text{st} \leq a \wedge (\bar{w}_1 \in T_1 \mid x_1 := w_1 \models A_1) \xrightarrow{\text{st}} (c_3 \mid s_3 \models A_3) \bullet \mathcal{R}(\text{st}, c_3) \} \cup \\ \{ \text{st}_1, c_2, s_2, \text{st}_2, c_3, s_3, A_3 \mid \\ \alpha \text{st}_1 \leq a_1 \wedge (\bar{w}_1 \in T_1 \mid x_1 := w_1 \models A_1) \xrightarrow{\text{st}_1} (c_2 \mid s_2 \models \text{Skip}) \wedge \\ \alpha \text{st}_2 \leq a_2 \wedge (\bar{w}_2 \in T_2 \mid x_2 := w_2 \models A_2) \xrightarrow{\text{st}_2} (c_3 \mid s_3 \models A_3) \\ \bullet \mathcal{N}^a(\mathcal{R}(\text{st}_1 \wedge \text{st}_2, c_2 \wedge c_3)) \}$$

[relational image, and definitions of ccat, *cstraces*, and *tcstraces*]

$$= \{ \text{st}, c_3, s_3, A_3 \mid \alpha \text{st} \leq a \wedge (\bar{w}_1 \in T_1 \wedge \bar{w}_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{\text{st}} (c_3 \mid s_3 \models A_3) \bullet \mathcal{R}(\text{st}, c_3) \} \cup \\ \{ \text{st}_1, c_2, s_2, \text{st}_2, c_3, s_3, A_3 \mid \\ \alpha \text{st}_1 \leq a_1 \wedge \\ (\bar{w}_1 \in T_1 \wedge \bar{w}_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{\text{st}_1} (c_2 \wedge \bar{w}_2 \in T_2 \mid s_2; x_2 := w_2 \models \text{Skip}) \wedge \\ \alpha \text{st}_2 \leq a_2 \wedge (c_2 \wedge \bar{w}_2 \in T_2 \mid s_2; x_2 := w_2 \models A_2) \xrightarrow{\text{st}_2} (c_3 \mid s_3 \models A_3) \\ \bullet \mathcal{N}^a(\mathcal{R}(\text{st}_1 \wedge \text{st}_2, c_2 \wedge c_3)) \}$$

[Propositions 1 and 2, $w_2 \neq w_1$, $x_2 \notin FV(A_1)$]

[by construction: $w_2 \notin SV(s_2, s_3)$, $w_2 \cap \alpha(\text{st}, \text{st}_1, \text{st}_2) = \emptyset$]

[by construction: $FV(c_2) \cap (\{w_2\} \cup SV(s_3)) = \emptyset$, $FV(c_2) \cap \alpha \text{st}_2 = \emptyset$]

[$FV(s_2) \cap FV(A_2) = \emptyset$, since $FV(s_2) \subseteq FV(s_1, A_1)$]

$$\begin{aligned}
&= \{ st, c_3, s_3, A_3 \mid \alpha st \leq a \wedge (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st} (c_3 \mid s_3 \models A_3) \bullet \mathcal{R}(st, c_3) \} \\
&\cup \\
&\{ st_1, c_2, s_2, st_2, c_3, s_3, A_3 \mid \\
&\quad \alpha st_1 \leq a_1 \wedge \\
&\quad (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st_1} (c_2 \wedge w_2 \in T_2 \mid s_2; x_2 := w_2 \models \text{Skip}) \wedge \\
&\quad \alpha st_2 \leq a_2 \wedge (c_2 \wedge w_2 \in T_2 \mid s_2; x_2 := w_2 \models A_2) \xrightarrow{st_2} (c_2 \wedge w_2 \in T_2 \wedge c_3 \mid s_3 \models A_3) \\
&\quad \bullet \mathcal{N}^a(\mathcal{R}(st_1 \hat{\ } st_2, c_2 \wedge w_2 \in T_2 \wedge c_3)) \}
\end{aligned}$$

[Proposition 3]

$$\begin{aligned}
&= \{ st, c_3, s_3, A_3 \mid \alpha st \leq a \wedge (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st} (c_3 \mid s_3 \models A_3) \bullet \mathcal{R}(st, c_3) \} \\
&\cup \\
&\{ st, st_1, c_2, s_2, st_2, c_3, s_3, A_3 \mid \\
&\quad st = st_1 \hat{\ } st_2 \wedge \alpha st_1 \leq a_1 \wedge \alpha st_2 \leq a_2 \wedge \\
&\quad (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st_1} (c_2 \mid s_2 \models \text{Skip}) \wedge \\
&\quad (c_2 \mid s_2 \models A_2) \xrightarrow{st_2} (c_3 \mid s_3 \models A_3) \\
&\quad \bullet \mathcal{N}^a(\mathcal{R}(st, c_3)) \}
\end{aligned}$$

[property of sets and Proposition 3]

$$\begin{aligned}
&= \{ st, c_3, s_3, A_3 \mid \alpha st \leq a \wedge (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st} (c_3 \mid s_3 \models A_3) \\
&\quad \bullet \mathcal{N}^a(\mathcal{R}(st, c_3)) \} \cup \\
&\{ st, st_1, c_2, s_2, st_2, c_3, s_3, A_3 \mid \\
&\quad st = st_1 \hat{\ } st_2 \wedge \alpha st_1 \leq a_1 \wedge \alpha st_2 \leq a_2 \wedge \\
&\quad (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st_1} (c_2 \mid s_2 \models \text{Skip}) \wedge \\
&\quad (c_2 \mid s_2 \models A_2) \xrightarrow{st_2} (c_3 \mid s_3 \models A_3) \\
&\quad \bullet \mathcal{N}^a(\mathcal{R}(st, c_3)) \}
\end{aligned}$$

$[\mathcal{N}^a(\mathcal{R}(st, c_3)) = \mathcal{R}(st, c_3) \text{ because } \alpha st \leq a]$

$$\begin{aligned}
&= \{ st, c_3, s_3, A_3, st_1, c_2, s_2, st_2 \mid \\
&\quad (\alpha st \leq a \wedge (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st} (c_3 \mid s_3 \models A_3)) \\
&\quad \vee \\
&\quad \left(\begin{array}{l} st = st_1 \hat{\ } st_2 \wedge \alpha st_1 \leq a_1 \wedge \alpha st_2 \leq a_2 \wedge \\ (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st_1} (c_2 \mid s_2 \models \text{Skip}) \wedge \\ (c_2 \mid s_2 \models A_2) \xrightarrow{st_2} (c_3 \mid s_3 \models A_3) \\ \bullet \mathcal{N}^a(\mathcal{R}(st, c_3)) \end{array} \right)
\end{aligned}$$

[property of sets]

$$\begin{aligned}
&= \{ st, c_3, s_3, A_3, st_1, c_2, s_2, st_2 \mid \alpha st \leq a \wedge \\
&\quad \left(\begin{array}{l} (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st} (c_3 \mid s_3 \models A_3) \\ \vee \\ \left(\begin{array}{l} st = st_1 \hat{\ } st_2 \wedge (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1) \xrightarrow{st_1} (c_2 \mid s_2 \models \text{Skip}) \wedge \\ (c_2 \mid s_2 \models A_2) \xrightarrow{st_2} (c_3 \mid s_3 \models A_3) \end{array} \right) \end{array} \right) \\
&\quad \bullet \mathcal{R}(st, c_3) \}
\end{aligned}$$

[property of alphabets and $\mathcal{N}^a(cst) = cst$ provided $\alpha cst = a$]

$$\begin{aligned}
&= \{ st, c_3, s_3, A_3 \mid \alpha st \leq a \wedge (w_1 \in T_1 \wedge w_2 \in T_2 \mid x_1, x_2 := w_1, w_2 \models A_1; A_2) \xrightarrow{st} (c_3 \mid s_3 \models A_3) \\
&\quad \bullet \mathcal{R}(st, c_3) \}
\end{aligned}$$

[Proposition 5]

$$= cstraces^a(w_1 \in T_1 \wedge w_2 \in T_2, (x_1, x_2 := w_1, w_2), A_1; A_2)$$

[definition of $cstraces$]

$$= cstraces^a(P_1; P_2)$$

[definition of process sequence and $cstraces$]

Case: hiding. We can assume, without loss of generality, that

$$P = \text{begin state } [x : T] \bullet A \text{ end}$$

$$\begin{aligned}
& \text{select}S^a(P_1 \setminus cs) \\
&= \mathcal{N}^a (\downarrow \text{select}S^a(P) \downarrow \upharpoonright_C cs) && \text{[definition of } \text{select}S \text{]} \\
&\subseteq \mathcal{N}^a (\downarrow \text{cstraces}^a(P) \downarrow \upharpoonright_C cs) && \text{[induction hypothesis]} \\
&= \mathcal{N}^a (\downarrow \text{cstraces}^a(w \in T, x := w, A) \downarrow \upharpoonright_C cs) && \text{[definition of } \text{cstraces} \text{]} \\
&= \{ \text{cst} : \text{cstraces}^a(w \in T, x := w, A) \bullet \mathcal{N}^a(\text{cst} \upharpoonright_C cs) \} && \text{[relational image]} \\
&= \{ \text{st}, c_2, s_2, A_2 \mid \alpha \text{st} \leq a \wedge (w \in T \mid x := w \models A) \xrightarrow{\text{st}} (c_2 \mid s_2 \models A_2) \bullet \mathcal{N}^a(\mathcal{R}(\text{st}, c_2) \upharpoonright_C cs) \} && \text{[definition of } \text{cstraces}^a \text{]} \\
&= \{ \text{st}, c_3, s_3, A_3 \mid \alpha \text{st} \leq a \wedge (w \in T \mid x := w \models A \setminus cs) \xrightarrow{\text{st}} (c_3 \mid s_3 \models A_3) \bullet \mathcal{R}(\text{st}, c_3) \} && \text{[Proposition 7]} \\
&= \text{cstraces}^a(P_1 \setminus cs) && \text{[definition of } \text{cstraces}^a \text{]}
\end{aligned}$$

□