



HAL
open science

The hitchhiker's guide to decidability and complexity of equivalence properties in security protocols (technical report)

Vincent Cheval, Steve Kremer, Itsaka Rakotonirina

► To cite this version:

Vincent Cheval, Steve Kremer, Itsaka Rakotonirina. The hitchhiker's guide to decidability and complexity of equivalence properties in security protocols (technical report). [Technical Report] Inria Nancy Grand-Est. 2020. hal-02501577v1

HAL Id: hal-02501577

<https://hal.science/hal-02501577v1>

Submitted on 9 Mar 2020 (v1), last revised 25 May 2020 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The hitchhiker’s guide to decidability and complexity of equivalence properties in security protocols (technical report¹)

Vincent Cheval, Steve Kremer, Itsaka Rakotonirina

Inria Nancy Grand-Est and LORIA

ABSTRACT

Privacy-preserving security properties in cryptographic protocols are typically modelled by observational equivalences in process calculi such as the applied pi-calculus. We survey decidability and complexity results for the automated verification of such equivalences, casting existing results in a common framework which allows for a precise comparison. This unified view, beyond providing a clearer insight on the current state of the art, allowed us to identify some variations in the statements of the decision problems—sometimes resulting into different complexity results. Additionally, we prove a couple of novel or strengthened results.

Keywords. Formal verification · cryptographic protocols · complexity.

1 INTRODUCTION

Symbolic verification techniques for security protocols can be traced back to the seminal work of Dolev and Yao [DY81]. Today, after more than 30 years of active research in this field, efficient and mature tools exist, e.g. `PROVERIF` [Bla16] and `TAMARIN` [SMCB13] to only name the most prominent ones. These tools are able to automatically verify full fledged models of widely deployed protocols and standards, such as TLS [BBK17, CHH⁺17], Signal [KBB17, CGCG⁺18], the upcoming 5G standard [BDH⁺18], or deployed multi-factor authentication protocols [JK18]. We argue that the development of such efficient tools has been possible due to a large amount of more theoretical work that focuses on understanding the precise limits of decidability and the computational complexity of particular protocol classes [DEK82, DLMS99, RT03, DLM04, CC05, KKNS14].

The abovementioned results extensively cover verification for the class of *reachability* properties. Such properties are indeed sufficient to verify authentication properties and various flavors of confidentiality, even in complex scenarios with different kinds of compromise [BC14]. Another class of properties are *indistinguishability* properties. These properties express that an adversary cannot distinguish two situations and is conveniently modelled as an *observational equivalence* in a cryptographic process calculus, such as the applied pi calculus. Such equivalences can indeed be used to model strong flavors of secrecy, in terms of non-interference or as a “real-or-random” experiment. Equivalences are also the tool of choice to model many other privacy-preserving properties. Such properties include anonymity [AF04], unlinkability properties [ACRR10, FHMS19], as well as vote privacy [DKR09]

¹This is the technical report of the survey [CKR20a]. It contains some proofs of claims that do not follow directly from the cited references, results that are considered folklore although not published, or are simply novel.

to give a few examples. Equivalence properties are inherently more complex than reachability properties, and both the theoretical understanding and tool support are more recent and more brittle. This state of affairs triggered a large amount of recent works to increase our theoretical understanding and improve tool support.

In this paper we give an extensive overview of decidability and complexity results for several process equivalences. In particular, in this survey we give a unified view, allowing us to highlight subtle differences in the definitions of the decision problems across the literature (such as whether the term theory is part of the input or not) as well as the protocol models. Typically, models may vary in whether they allow for a bounded or unbounded number of sessions, the support of cryptographic primitives, whether they support else branches (i.e. disequality tests, rather than only equality tests), and various restrictions on non-determinism. All the results are summarised in Table 1, and we identify several open questions.

2 MODEL

In this section we present the symbolic model of security protocols we consider, the applied pi-calculus [ABF17], rooted in the seminal work of Dolev and Yao [DY81]. Since the models used by the works we survey often differ in their presentation, we use a middleground, custom model allowing for expressing the cited theorems with minimal tweaking of their original statements. We assume the reader familiar with the theory of rewriting.

Cryptographic primitives As usual in symbolic protocol analysis we take an abstract view of cryptography and model the messages exchanged during the protocol as *terms* built over a set of function symbols each with a given arity called a *signature*. Terms are then either atomic values or function symbols applied to other terms, respecting the function’s arity. Atomic values are either *constants*, *names*, or *variables*. Constants, sometimes referred as public names, model public values such as agent identities or protocol tags. Names, sometimes explicitly called private names, model fresh secret values, such as keys or nonces, and are a priori unknown to the adversary. As usual variables express bound values and serve as domain for substitutions. We assume an infinite set of constants Σ_0 , names \mathcal{N} and variables \mathcal{X} and write $\mathcal{T}(\Sigma, A)$ the set of terms built from the signature Σ and atomic values of A .

Example 2.1. A signature Σ for symmetric encryption and pairs is usually written as follows

$$\Sigma = \{\text{senc}/2, \text{sdec}/2, \langle, \rangle/2, \text{fst}/1, \text{snd}/1\}.$$

For example, the encryption of a plaintext m with a key k would be modelled by the term $\text{senc}(m, k)$. To include a randomness nonce r , we can encrypt a pair which gives the term $\text{senc}(\langle m, r \rangle, k)$. \triangle

The functional properties of the symbols are modelled by an *equational theory*. In this work we restrict ourselves to equational theories that can be oriented into a *convergent rewriting system*. This also implies that any term t has a unique normal form $t\downarrow$.

Example 2.2. The rewrite rules

$$\text{sdec}(\text{senc}(x, y), y) \rightarrow x \quad \text{fst}(\langle x, y \rangle) \rightarrow x \quad \text{snd}(\langle x, y \rangle) \rightarrow y$$

define the behaviour of the pairs and the encryption scheme. Typically one can decrypt (apply sdec) a ciphertext $\text{senc}(x, y)$ with the corresponding key y to recover the plaintext x . This behaviour is idealised by the absence of other rules for senc and sdec , modelling an assumption that no information can be extracted from a ciphertext except by possessing the decryption key. \triangle

In this survey we call a *theory* the set of non-constant function symbols together with a rewriting system. They can express a broad range of other cryptographic primitives, like the following ones that will be used in this survey:

- *randomised symmetric encryption*, adding an explicit argument for a randomness nonce. It is defined by $\Sigma = \{\text{rsenc}/3, \text{rsdec}/2\}$ and $\text{rsdec}(\text{rsenc}(m, r, k), k) \rightarrow m$.
- *randomised asymmetric encryption*, which is its analogue with public-key mechanisms: $\Sigma = \{\text{pk}/1, \text{raenc}/3, \text{radec}/2\}$ and $\text{radec}(\text{raenc}(m, r, \text{pk}(k)), k) \rightarrow m$.
- *digital signature*, with a verification mechanism that recovers the signed message: $\Sigma = \{\text{pk}/1, \text{sign}/3, \text{verify}/2\}$ and $\text{verify}(\text{sign}(m, r, k), \text{pk}(k)) \rightarrow m$.
- *one-way hash*, simply using a function symbol of positive arity, e.g. $\Sigma = \{h/1\}$. One-wayness is modelled by an absence of rewrite rules involving h , in which case we say that h is *free*.

Two classes of theories are particularly important for our results. The first is the class of *subterm convergent* theories [AC06, Bau07, BAF08, CKR18a, CDK09, CBC11], defined by a syntactic criterion on rewriting rules $\ell \rightarrow r$ requiring that r is either a strict subterm of ℓ or a ground term in normal form. The second is the class of *constructor-destructor* theories [BAF08, CCLD11, CKR18a], partitioning function symbols into constructor (used to build terms) and destructors (only used in rewrite rules). In constructor-destructor theories any rewrite rule $\ell \rightarrow r$ is such that $\ell = d(t_1, \dots, t_n)$ where d is a destructor and t_1, \dots, t_n, r do not contain any destructor. Moreover, we assume a *message* predicate $\text{msg}(t)$ which holds if $u\downarrow$ does not contain any destructor symbol for all subterms u of t , i.e., all destructor applications in t succeeded yielding a valid message. This predicate is used to restrict to protocols that only send and accept such well-formed messages. All theories above are subterm convergent and constructor destructor.

Protocols Protocols are defined using *processes* in the applied pi calculus. Their syntax is defined by the following grammar:

$P, Q ::= 0$	(null process)
$\text{if } u = v \text{ then } P \text{ else } Q$	(conditional)
$u(x).P$	(input)
$\bar{u}\langle v \rangle.P$	(output)
$P \mid Q$	(parallel)

where u, v are terms and x a variable. Intuitively the 0 models a terminated process, a conditional if $u = v$ then P else Q executes either P or Q depending on whether the terms $u\downarrow$ and $v\downarrow$ are equal, and $P \mid Q$ models two processes executed concurrently. The constructs $c(x).P$ and $\bar{c}\langle u \rangle.P$ model, respectively, inputs and outputs on a communication channel c . When the channel c is known to the attacker, e.g. when it is a constant, executing an output on c adds it to the adversary's knowledge and inputs on c are fetched from the adversary possibly forwarding a previously stored message, or computing a new message from previous outputs. Otherwise the communication is performed silently without adversarial interferences. To model an unbounded number of a protocol session we also add the two constructs

$P, Q ::= \text{new } k.P$	(new name)
$!P$	(replication)

The replication $!P$ models an unbounded number of parallel copies of P , and $\text{new } k.P$ creates a fresh name k unknown to the attacker; in particular $!\text{new } k.P$ models an unbounded number of sessions, each with a different fresh key. The fragment of the calculus without replication is referred as *finite* or *bounded*. Another notable subclass is the original pi-calculus [MPW92], referred as the *pure* fragment, that can be retrieved with the empty theory (only names, constants and an empty rewrite system).

Attacker's knowledge We model the attacker's observations recorded when spying on the communication network by a *frame*. A frame is a substitution of the form

$$\Phi = \{ax_1 \mapsto t_1, \dots, ax_n \mapsto t_n\}$$

where t_i are the outputs performed during the execution of the protocol and $ax_i \in \mathcal{AX}$, with \mathcal{AX} a set of special variables called *axioms* that serve as handles to the adversary for building new terms. These terms t_i enable adversarial deductions as they aggregate: for example after observing a ciphertext and the decryption key, the attacker can also obtain the plaintext by decrypting. Formally we say that one can *deduce* all terms $\xi\Phi\downarrow$ where $\xi \in \mathcal{T}(\Sigma, \Sigma_0 \cup \text{dom}(\Phi))$ is called a *recipe*. A recipe models a computation of the adversary: the fact that it cannot contain names models that they are assumed unknown to her. They naturally only remain unknown while they are not revealed in the frame themselves; for example in

$$\Phi = \{ax_1 \mapsto \text{senc}(t, k), ax_2 \mapsto k\}$$

even if deducing the term t requires to decrypt $ax_1\Phi$ with the key k (which is not allowed to occur directly in the recipe), this is possible by using $\xi = \text{sdec}(ax_1, ax_2)$. We refer to the following decision problem as **DEDUCIBILITY**:

INPUT: a theory, a frame Φ , a term t

QUESTION: Does there exist a recipe ξ such that $\xi\Phi\downarrow = t\downarrow$?

Semantics in an adversarial environment The behaviour of processes is formalised by an operational semantics. The detailed presentation differs from one work to another [CCD13, ABF17, CKR18a, CKR19] and we choose a formalism that permits to state all theorems with minimal changes in the proofs. The semantics

operates on *extended processes* (\mathcal{P}, Φ) where \mathcal{P} is a multiset of processes modelling the state of the processes currently executed in parallel, and Φ is the frame indicating the outputs the attacker has recorded during the execution. It takes the form of a labelled transition relation $\xrightarrow{\alpha}$ whose label α is called an *action*, which is either

- a public input action $\xi_c(\xi_t)$ where ξ_c (resp. ξ_t) is a recipe for the input's channel (resp. of the term to be input);
- a public output action $\overline{\xi_c}\langle ax_i \rangle$ where ξ_c is a recipe for the output's channel, and the underlying output term is added to the frame under axiom ax_i ;
- an unobservable action τ which represents an internal action, such as the evaluation of a conditional or a communication on a private channel.

This is formalised in Figure 1. When defining security against an active attacker we quantify over all such transitions which means we consider all possible executions in an active adversarial environment. Thus even the bounded fragment yields an infinite transition system if the theory contains a non-constant function symbol (as this allows to build an unbounded number of messages).

Variations across the literature There are several modelling variations of this semantics. The most important one is when the theory is constructor-destructor. For this class of theories, in this survey, we always refer to an altered semantics that intuitively requires that all destructor operations succeed for a transition to be applied [CCD15b, CCD15a, CKR18a]. Formally:

- the communication rules (IN), (OUT), (COMM) are only applicable when all terms $\xi_c\Phi, \xi_t\Phi, u, u', v$ verify the predicate msg . For instance no transitions are possible from the process

$$\overline{c}\langle \text{sdec}(a, b) \rangle.P$$

with a, b, c constants because $\text{sdec}(a, b)$ is not a message.

- the conditional rule (TEST) executes its negative branch when a destructor fails, i.e. with the notations of Figure 1, $R = P$ if $\text{msg}(u)$, $\text{msg}(v)$ and $u \not\downarrow = v \downarrow$. In particular, although this may seem counterintuitive at first sight, the process

$$\text{if } \text{sdec}(a, b) = \text{sdec}(a, b) \text{ then } P \text{ else } Q$$

reduces to Q by Rule (TEST) in this semantics.

In the examples above with sdec , this constructor-destructor semantics models an assumption that the encryption scheme has enough structure to detect decryption failure, and that the protocol only proceeds with valid messages.

Besides, as noted in [BCK20], synchronous communications between parallel processes (Rule (COMM)) is also managed differently from one work to another. In the original semantics [ABF17] of the applied pi-calculus, called the *classical semantics* in [BCK20], communications on a same public channel between parallel processes can either be executed silently without adversarial interference (i.e. using (COMM)) or be routed through the attacker (i.e. using a sequence of (OUT) and (IN)). This is also the semantics used in the popular PROVERIF tool [BAF08]. On the contrary, the semantics defined in Figure 1 only allows applications of Rule (COMM) when the channel is unknown to the adversary, modelling an attacker that continuously eavesdrops on the network (rather than

an attacker that solely has the capability to do so). This is called the *private semantics* in [BCK20]. The private semantics is actually used in tools such as TAMARIN [SMCB13] and AKISS [CCCK16] and also in a few other works we survey [CCD15b, CCD15a, CKR19].

While both semantics are equivalent when it comes to reachability properties, they surprisingly happen to be incomparable for equivalence properties [BCK20]. All the complexity results of this paper are with respect to the private semantics. Although we did not expand on studying all the variations of complexity induced by using different semantics, most of the analyses presented in this survey are robust to these changes. Indeed, all complexity results for the bounded fragment hold for both semantics. In the unbounded case, only the private semantics has been considered in the underlying models [CCD15a, CCD15b].

3 COMPLEXITY FOR A PASSIVE ATTACKER

3.1 Static equivalence

Some security properties against a passive attacker, i.e. a simple eavesdropper, can then be modelled as an observational equivalence of two frames: intuitively no equality test can be used to distinguish them. For example, in a protocol that outputs a sequence of messages t_1, \dots, t_n , the “real-or-random” confidentiality of a key k can be modelled as the equivalence of

$$\Phi = \{ax_1 \mapsto t_1, \dots, ax_n \mapsto t_n, ax \mapsto k\}$$

$$\Psi = \{ax_1 \mapsto t_1, \dots, ax_n \mapsto t_n, ax \mapsto k'\}$$

where k' is a fresh name. More formally, two frames Φ, Ψ with same domain are *statically equivalent* when for all recipes ξ_1, ξ_2 ,

$$\xi_1\Phi \downarrow = \xi_2\Phi \downarrow \iff \xi_1\Psi \downarrow = \xi_2\Psi \downarrow .$$

In constructor destructor theories we also require that $\text{msg}(\xi_1\Phi)$ iff $\text{msg}(\xi_1\Psi)$, modelling an assumption that the adversary can observe destructor failures.

Example 3.1. If k, k' are names, $\Phi = \{ax \mapsto k\}$ and $\Psi = \{ax \mapsto k'\}$ are statically equivalent, capturing the intuition that random keys cannot be distinguished. Similarly, the frames $\Phi = \{ax \mapsto k\}$ and $\Psi = \{ax \mapsto \text{senc}(t, k')\}$ are statically equivalent for any term t , modelling that encryption is indistinguishable from a random string. However, for the constant 0,

$$\Phi = \{ax_1 \mapsto \text{senc}(0, k), ax_2 \mapsto k\}$$

$$\Psi = \{ax_1 \mapsto \text{senc}(0, k), ax_2 \mapsto k'\}$$

are not statically equivalent since $\xi_1 = \text{sdec}(ax_1, ax_2)$ and $\xi_2 = 0$ are equal in Φ but not in Ψ . \triangle

3.2 Complexity results

We survey the decidability and complexity of the following decision problem referred as STATEQ:

INPUT: A theory, two frames of same domain.

QUESTION: Are the frames statically equivalent for this theory?

General case. As rewriting is Turing-complete, unsurprisingly static equivalence is undecidable in general for convergent rewriting systems [AC06]. It is also proved in [AC06] that DEDUCIBILITY

$(\{\{u(x).P\}\} \cup \mathcal{P}, \Phi) \xrightarrow{\xi_c(\xi_t)} (\{\{P\{x \mapsto \xi_t \Phi\}\}\} \cup \mathcal{P}, \Phi)$	if $\xi_c \Phi \downarrow = u \downarrow$	(IN)
$(\{\{\bar{u}(v).P\}\} \cup \mathcal{P}, \Phi) \xrightarrow{\xi_c(ax)} (\{\{P\}\} \cup \mathcal{P}, \Phi \cup \{ax \mapsto v \downarrow\})$	if $\xi_c \Phi \downarrow = u \downarrow$ and $ax \in \mathcal{AX} \setminus \text{dom}(\Phi)$	(OUT)
$(\{\{\bar{u}(v).P, u'(x).Q\}\} \cup \mathcal{P}, \Phi) \xrightarrow{\tau} (\{\{P, Q\{x \mapsto v\}\}\} \cup \mathcal{P}, \Phi)$	if $u \downarrow = u' \downarrow$ and u not deducible from Φ	(COMM)
$(\{\{\text{if } u = v \text{ then } P \text{ else } Q\}\} \cup \mathcal{P}, \Phi) \xrightarrow{\tau} (\{\{R\}\} \cup \mathcal{P}, \Phi)$	where $R = P$ if $u \downarrow = v \downarrow$ and $R = Q$ otherwise	(TEST)
$(\{\{\text{new } k.P\}\} \cup \mathcal{P}, \Phi) \xrightarrow{\tau} (\{\{P\{k \mapsto k'\}\}\} \cup \mathcal{P}, \Phi)$	if k' is a fresh name	(NEW)
$(\{\{P \mid Q\}\} \cup \mathcal{P}, \Phi) \xrightarrow{\tau} (\{\{P, Q\}\} \cup \mathcal{P}, \Phi)$		(PAR)
$(\{\{!P\}\} \cup \mathcal{P}, \Phi) \xrightarrow{\tau} (\{\{!P, P\}\} \cup \mathcal{P}, \Phi)$		(REPL)

Figure 1: Operational semantics of the applied pi-calculus

reduces to STAT_{EQ} . As a consequence, the results of [ANR07] imply that static equivalence is also undecidable for so-called *optimally-reducing* rewrite systems, a subclass of rewrite systems that have the finite-variant property [CCCK16].

Subterm convergent theories. Historically, the complexity of static equivalence has only been considered for *fixed* theories [AC06, Bau07], that is, the theory was not part of the input of the problem and its size was seen as a constant in the complexity analysis. This was consistent with most formalisms and verification tools at the time, which would not allow for user-defined theories and only consider a fixed set of cryptographic primitives, such as in the spi-calculus for example [AG99]. For example:

THEOREM 3.1 ([AC06]). For all fixed subterm convergent theories STAT_{EQ} is PTIME.

A generic PTIME-completeness result would make no sense when the theory is not part of the input, since the complexity may depend of it. Typically:

THEOREM 3.2 ([CKR18a]). In the pure pi-calculus (i.e. with an empty theory) STAT_{EQ} is LOGSPACE.

However the PTIME bound is optimal in the following sense:

THEOREM 3.3. For all fixed theories containing symmetric encryption, STAT_{EQ} is PTIME-hard.

Proof sketch. We proceed by reduction from HORNSAT . Let X be the set of variables of a Horn formula $\varphi = C_1 \wedge \dots \wedge C_n$, and k_x be names for all $x \in X \cup \{\perp\}$. Then to each clause $C_i = x_1, \dots, x_n \Rightarrow x, x \in X \cup \{\perp\}$ we associate the term

$$t_{C_i} = \text{senc}(\dots \text{senc}(\text{senc}(k_x, k_{x_1}), k_{x_2}), \dots, k_{x_n}).$$

Then k_{\perp} is deducible from terms t_{C_1}, \dots, t_{C_n} iff the formula φ is unsatisfiable. In particular given two constants 0,1, and $\Phi = \{ax_1 \mapsto t_{C_1}, \dots, ax_n \mapsto t_{C_n}\}$, then the frames

$$\Phi \cup \{ax \mapsto \text{senc}(0, k_{\perp})\} \quad \text{and} \quad \Phi \cup \{ax \mapsto \text{senc}(1, k_{\perp})\}$$

are statically equivalent iff φ is satisfiable. \square

However automated tools have improved since then and some provers like KISS [CDK09], YAPA [BCD13] or FAST [CBC11] are able to handle user-defined theories. It is therefore interesting today to account for the size of the theory in the complexity analysis:

THEOREM 3.4 ([CKR18a]). STAT_{EQ} is coNP-complete for subterm convergent theories.

Proof sketch. We sketch the reduction from SAT presented in [CKR18a]. We consider two constants 0 and 1, function symbols f, g of arity 2, and the two frames

$$\begin{aligned} \Phi &= \{ax_0 \mapsto f(0, k), ax_1 \mapsto f(1, k)\} \\ \Psi &= \{ax_0 \mapsto g(0, k), ax_1 \mapsto g(1, k)\} \end{aligned}$$

for some name k . Interpreting 0 and 1 as the booleans false and true, Φ and Ψ point to terms that can be seen as booleans *but* that can only be accessed by reference through the axioms ax_0, ax_1 . For example, since k is a name the only recipe permitting to deduce $f(0, k)$ is ax_0 in Φ . Given a SAT formula φ of variables x_1, \dots, x_n , we then add an other symbol eval of arity n and rewrite rules so that the following points are equivalent for all valuations $v : \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ of φ :

$$(1) v \text{ falsifies } \varphi \quad (2) \text{eval}(g(v(x_1), k), \dots, g(v(x_n), k)) \rightarrow 0$$

Details can be found in [CKR18a]. If we add the rule

$$\text{eval}(f(y_1, z), \dots, f(y_n, z)) \rightarrow 0$$

we eventually have that φ is satisfiable iff there exists a valuation v such that $t_v \Psi \neq 0$ where $t_v = \text{eval}(ax_{v(x_1)}, \dots, ax_{v(x_n)})$, iff Φ and Ψ are not statically equivalent. \square

Beyond subterm convergence Although we are not aware of complexity results for the decision of static equivalence for classes larger than subterm theories, there exist decidability results. Some of the abovementioned tools, like KISS and YAPA, can actually handle most convergent rewriting system; but they naturally fail to terminate in general by undecidability of the problem. However it is proved for example in [CDK09] that the termination of KISS is guaranteed for theories modelling blind signatures or trapdoor commitment schemes (that are typically not subterm).

4 COMPLEXITY FOR AN ACTIVE ATTACKER

In this section we survey the decidability and complexity of equivalence relations characterising security against active attackers.

4.1 Equivalences

We expect security protocols to provide privacy-type guarantees against attackers that actively engage with the protocol. This can be modelled by behavioural equivalences, defining security as the indistinguishability of two instances of the protocol that differ on a privacy-sensitive attribute such as a secret key, an identity, or the agent executing a given session. There exist several candidate equivalences for modelling this notion of indistinguishability. We study two of them in this survey and refer to [CCD13] for a more detailed overview and comparison with other equivalences.

Trace equivalence One classical example of such behavioural equivalence is *trace equivalence*. Referring to the operational semantics mentioned in Figure 1, we call a *trace* of a process P a sequence of transition steps from P in this semantics, i.e.

$$(\llbracket P \rrbracket, \emptyset) = A_0 \xrightarrow{\alpha_1} A_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} A_n \quad \text{written } A_0 \xrightarrow{\alpha_1 \dots \alpha_n} A_n$$

for extended processes A_1, \dots, A_n . Given such a trace t , we write $actions(t) = \alpha_1 \dots \alpha_n$ the sequence of actions taken by the trace, and $\Phi(t)$ the frame of A_n , that is, the knowledge of the attacker at the end of the trace. In particular t and t' are said equivalent, written $t \sim t'$ here, when $actions(t)$ and $actions(t')$ are identical after erasure of τ actions and $\Phi(t)$ and $\Phi(t')$ are statically equivalent.

Two processes P_0 and P_1 are said trace equivalent when for all traces t of P_i , $i \in \{0, 1\}$, there exists a trace t' of P_{1-i} such that $t \sim t'$. Trace equivalence has been studied intensively for the automation of security proofs [CCLD11, CCD13, ACK16, CKR18a] and has received a strong tool support [Che14, CCCK16, CGLM17, CKR18b, CDD18]. We refer to its decision problem as `TRACEEQ`:

INPUT: A theory, two processes.

QUESTION: Are the two processes trace equivalent?

Labelled bisimilarity Some other automated tools aim at proving more fine-grained equivalence, like *observational equivalence* for `PROVERIF` [BAF08, CB13] for example. There exist several flavours of more operational bisimulation-based properties but the one that is usually considered in security-protocol analysis is *labelled bisimilarity* because it coincides with observational equivalence in the applied pi-calculus [ABF17]. Formally it is an early, weak bisimulation that additionally requires static equivalence at each step; that is, it is the largest symmetric binary relation \approx on processes such that $A \approx B$ implies

- the frames of A and B are statically equivalent
- for all transitions $A \xrightarrow{\alpha} A'$, there exists $B \xrightarrow{\tau \dots \tau \cdot \alpha \cdot \tau \dots \tau} B'$ such that $A' \approx B'$.

We refer to the following problem as `BISIM`:

INPUT: A theory, two processes.

QUESTION: Are the two processes labelled bisimilar?

4.2 Notorious fragments

In addition to the assumptions on the rewriting system (e.g. sub-term convergence as in Section 3), there are several common restrictions made on the processes to obtain decidability.

Conditionals and patterns A typical restriction on conditionals is the class of *positive* processes that only contain trivial else branches [Bau07, CCD13, CKR18a]. For succinctness we write

$$[u = v]P \quad \text{instead of} \quad \text{if } u = v \text{ then } P \text{ else } 0.$$

When the rewrite system is constructor-destructor, some conditionals may also be encoded within inputs [CCD15a, CCD15b], using a notation $u(v).P$ where v is a term without destructors (but may contain variables) that is called a *pattern* in this survey. In terms of semantics, the transition rule (IN) is generalised to:

$$(\llbracket u(v).P \rrbracket \cup \mathcal{P}, \Phi) \xrightarrow{\xi u(\xi v)} (\llbracket P\sigma \rrbracket \cup \mathcal{P}, \Phi) \quad (\text{P-IN})$$

if $\xi v \Phi \downarrow = v\sigma \downarrow$, as well as the usual conditions $\text{msg}(u)$, $\text{msg}(\xi u \Phi)$, $\text{msg}(\xi v \Phi)$, $\xi u \Phi \downarrow = u \downarrow$. For example a process $c(\text{senc}(x, k)).P$ only reads inputs that are terms t encrypted with the key k , and x will then be bound to t in P . In this paper, to ensure that protocols can be effectively implemented we require that

- ▶ It is possible to test with a sequence of positive conditionals that a term t matches the pattern v .

That is, there exist terms $t_1, \dots, t_n, t'_1, \dots, t'_n$ (possibly containing a variable x) such that for all ground terms t , t is an instance of v iff for all $i \in \llbracket 1, n \rrbracket$, $t_i \{x \mapsto t\} \downarrow = t'_i \{x \mapsto t\} \downarrow$. This excludes patterns like $\langle \text{rsenc}(x, y, z), \text{rsenc}(x', y, z') \rangle$ that would accept any pair of ciphertexts encrypted using the same randomness.

- ▶ All free variables of v effectively appearing in the rest of the process can be extracted by applying destructors to v .

That is, for all variables x of v that are free (i.e. are not bound by a previous input) and appear in P , there exists a term context C without free variables such that $C[v] \downarrow = x$. This excludes for example patterns $h(x)$ where h is a free function symbol: given an input term $h(t)$, the one-wayness of h prevents from retrieving t . The assumption that C does not contain free variables excludes, for example, patterns $\text{senc}(0, y)$ that would accept the constant 0 encrypted by any key. On the contrary, a pattern $v = \text{rsenc}(x, y, k)$ is valid if $k \in \mathcal{N}$ and the variable y does not appear in P .

All in all, we define the *patterned* fragment to be the class of processes without conditionals but using pattern inputs, and where outputs do not contain destructor symbols; it is a subset of the positive fragment.

Ping pong protocols These protocols [CCD15b, DY81, HS03] consist of an unbounded number of parallel processes receiving one message and sending a reply. Although the precise formalisms may differ from one work to another, the mechanisms at stake are essentially captured by processes $P = !P_1 \mid \dots \mid !P_n$ where each P_i can be written under the form

$$P_i = c_i(x). [u_1^i = v_1^i] \cdot \dots [u_{n_i}^i = v_{n_i}^i] \text{new } k_1 \dots \text{new } k_{r_i}. \overline{c_i} \langle w_i \rangle$$

In particular ping-pong protocols are positive.

Simple processes. A common middleground in terms of expressivity and decidability is the class of simple processes, for example studied in [CCD13, CCD15a]. Intuitively, they consist of a sequence of parallel processes that operate each on a distinct, public channel—including replicated processes that generate dynamically a fresh channel for each copy. Formally they are of the form

$$P_1 \mid \dots \mid P_m \mid !^{ch} P_{m+1} \mid \dots \mid !^{ch} P_n. !^{ch} P = ! \text{new } c_P. \overline{c_P} \langle c_P \rangle. P$$

where each P_i does not contain parallel operators nor replications and uses a unique, distinct communication channel c_{P_i} , and

$$!^{ch} P = ! \text{new } c_P. \overline{c_P} \langle c_P \rangle. P.$$

Unlike ping pong protocols, each parallel process may input several messages and output messages that depend on several previous inputs. There exists a generalisation of simple processes called *determinate processes*, mentioned later in Section 6.

4.3 Complexity results: bounded fragment

The bounded fragment is a common restriction to study decidability, as removing replication bounds the length of traces. However, as the attacker still has an unbounded number of possibilities for generating inputs, the transition system still has infinite branching in general. Besides additional restrictions are necessary on the cryptographic primitives (at least because static equivalence is undecidable in general). For example:

THEOREM 4.1 ([CKR18a]). TRACEQ and BISIM are CONEXP for subterm convergent constructor destructor theories and bounded processes.

In a nutshell, the decision procedures use a dedicated constraint solving approach to show that, whenever trace equivalence is violated, there exists an attack trace whose attacker-input terms are at most of exponential size; in particular this shows non-equivalence to be decidable in NEXP . As before, we may also study the problem for fixed theories to investigate their influence on the complexity; typically with the empty theory:

THEOREM 4.2 ([CKR18a]). In the pure pi-calculus, TRACEQ (resp. BISIM) is Π_2 -complete (resp. PSPACE -complete) for bounded processes, and for bounded positive processes.

However, unlike static equivalence, fixing the theory does not make it possible to obtain a better bound than the general one:

THEOREM 4.3 ([CKR18a]). There exists a subterm convergent constructor destructor theory such that TRACEQ and BISIM are CONEXP -hard for bounded positive processes.

The theory in question [CKR18a] encodes binary trees and a couple of ad hoc functionalities. We show in Appendix A that, provided we discard the positivity requirement, it is possible to manage to the proof with a theory limited to symmetric encryption and pairs. This shows that the problem remains theoretically hard even with a minimal theory. Besides, in the case of trace equivalence, we also show that all abovementioned reductions can be done with only constants as channels (whereas [CKR18a] heavily relies on private communications, which may give the false intuition that they are necessary to obtain this high complexity).

4.4 Complexity results: unbounded fragment

Equivalence is undecidable in general since the calculus is Turing-complete even for simple theories. For example, Hüttel [Hüt03] shows that Minsky’s two counter machines can be simulated within the spi-calculus (and hence the applied pi-calculus with symmetric encryption only). It is not difficult to adapt the proof to a simulation using only a free symbol, i.e., a function symbol h of positive arity and an empty rewrite system. These two encodings can be performed within the *finite-control fragment*, typically not Turing-complete in the pure pi-calculus (i.e. without this free function symbol) [Dam97].

Ping pong protocols. While equivalence is undecidable for ping-pong protocols [CCD15b, HS03] some results exist under additional assumptions. For example [HS03] studies a problem that can be described in our model essentially as BISIM for ping-pong protocols with 2 participants or less (i.e. $n \leq 2$ in the definition). This is proved decidable under some model-specific assumptions which we do not detail here. We also mention a result for patterned ping-pong protocols (cf Section 4.2) without a limit on the number of participants [CCD15b]. Given a constructor-destructor theory, a ping-pong protocol P is said *deterministic* when each P_i (using the same notations as in the definition) can be written

$$P_i = c_i(u_i). \text{new } k_1 \dots \text{new } k_{r_i}. \overline{c_i} \langle v_i \rangle$$

with c_i a constant and u_1, \dots, u_n a family of patterns verifying:

- (1) *binding uniqueness*: for all i, u_i does not contain two different variables;
- (2) *pattern determinism*: for all $i \neq j$, if u_i and u_j are unifiable then $c_i \neq c_j$.

There is an additional syntactic restriction on the structures of u_i and v_i that is specific to the fixed theory considered in [CCD15b], containing randomised symmetric and asymmetric encryption and digital signature. The two terms u_i, v_i are defined by grammars essentially imposing that the subterms that serve as randomness (resp. keys) are indeed fresh nonces (resp. long-term keys), that is, they are names among k_1, \dots, k_{r_i} (resp. are of the form k or $\text{pk}(k)$ for some name $k \notin \{k_1, \dots, k_{r_i}\}$). We refer to [CCD15b] for details about this last assumption.

THEOREM 4.4 ([CCD15b]). For a theory limited to randomised symmetric and asymmetric encryption as well as digital signature, TRACEQ is decidable in primitive recursive time for deterministic ping-pong protocols.

Decidability is obtained by a reduction of the problem to the language equivalence of deterministic pushdown automata, which is decidable in primitive recursive time. A complexity lower bound for this problem is open (beyond the PTIME -hardness inherited from static equivalence, recall Theorem 3.3).

For simple processes We now study a decidability result for patterned simple processes [CCD15a]. It relies on three restrictions: the theory is limited to symmetric encryption and pairs, and the processes must be *type compliant* and *acyclic*. We give an intuition of the definition which is formalised in Appendix B. Type compliance relies on a type system to ensure that, whenever

two (subterms of) terms u, v appearing in the process are unifiable then, once input variables are bound to adversary-chosen terms, u and v have the same structure in terms of encryptions and pairs. On the other hand, acyclicity is a property of the *dependency graph* of the process. The vertices of this graph are the inputs and outputs of the process. There is an edge $a \rightarrow a'$ when it is necessary to execute a' before a to be able to perform some attacker actions.

Example 4.1. There are three kind of edges in a dependency graph. *Sequential dependency* is for actions following each other, for example in $\beta.\alpha.P$ there is an edge $\alpha \rightarrow \beta$. *Pattern* and *deduction dependencies* are for actions that allow the attacker to produce a term of a given pattern or deduce a subterm of an output message, respectively. For example in $\alpha.P \mid \beta.Q \mid \gamma.R$ with

$$\alpha = \bar{c}\langle \text{senc}(u, k) \rangle \quad \beta = d(\text{senc}(x, k)) \quad \gamma = \bar{e}\langle k \rangle$$

there is an edge $\alpha \rightarrow \beta$ because the term $\text{senc}(u, k)$ could be used as an input term for the pattern $\text{senc}(x, k)$. Also $\gamma \rightarrow \alpha$ because the term k output in γ can be used to deduce u from $\text{senc}(u, k)$ in α . Similarly note that there is a cyclic dependency in

$$\text{!}^{\text{ch}} \beta.\alpha \quad \text{with} \quad \alpha = \bar{c}\langle \text{senc}(u, k) \rangle \quad \beta = c(\text{senc}(x, k)).$$

We have $\alpha \rightarrow \beta$ by sequential dependency, but also $\beta \rightarrow \alpha$ by pattern dependency across the different copies of $\beta.\alpha$. \triangle

THEOREM 4.5 ([CCD15a]). For a theory limited to pairs and symmetric encryption, TRACEQ is coNEXP for patterned, simple, type-compliant, acyclic processes.

Proof. Given a trace we consider its so-called *execution graph*: its vertices are the actions of the trace and its edges mirror those of the dependency graph of the process. It is proved in [CCD15a] that when two patterned, simple, type-compliant, acyclic processes P and Q are not trace equivalent, there exists an attack trace, say, in P , whose execution graph D has these properties:

- (1) D is acyclic and $\text{depth}(D)$ (maximal length of a path of D) is polynomial in the size of P .
- (2) $\text{width}(D)$ (maximal number of outgoing edges from a vertex of D) is exponential in the size of P and of the type system.
- (3) $\text{nbroots}(D)$ (number of vertices of D that have no ingoing edges) is exponential in the size of P and of the type system.

From each root of D , the number of reachable vertices is at most the size of a tree of width $\text{width}(D)$ and of depth $\text{depth}(D)$, i.e. $\text{width}(D)^{\text{depth}(D)+1} - 1$. Hence the number of vertices of D is bounded by $\text{nbroots}(D) \cdot \text{width}(D)^{\text{depth}(D)+1}$ which is exponential in the size of P . Since the number of vertices of D is an upper bound on the number of sessions needed to execute the underlying trace, it suffices to prove the equivalence of P and Q for an exponential number number of sessions. This leads to an overall coNEXP procedure since trace equivalence of bounded, positive, simple processes is coNP for subterm theories (see Section 5). \square

Complexity was not the focus of [CCD15a] and the authors only claimed a triple exponential complexity for their procedure. Besides no lower bounds were investigated, but we proved that the problem was coNEXP -complete.

THEOREM 4.6. For the theory of pairs and symmetric encryption, TRACEQ is coNEXP -hard for patterned, simple, type-compliant, acyclic processes.

The reduction shares some similarities with the proof of coNEXP hardness for trace equivalence of bounded processes (see Theorem 4.3), compensating the more deterministic structure of simple processes by the use of replication. We give below an intuition of our construction, detailed in Appendix B.

Proof sketch. We proceed by reduction from SUCCINCT 3SAT . This is a common NEXP -complete problem that, intuitively, is the equivalent of 3SAT for formulas of exponential size represented succinctly by boolean circuits. Formally a formula φ with 2^m clauses and 2^n variables x_0, \dots, x_{2^n-1} is encoded by a circuit $\Gamma : \{0, 1\}^{m+2} \rightarrow \{0, 1\}^{n+1}$ in the following way. If $\varphi = \bigwedge_{i=0}^{2^m-1} \ell_i^1 \vee \ell_i^2 \vee \ell_i^3$ and $0 \leq i \leq 2^m - 1$ and $0 \leq j \leq 2$, we let x_k the variable of the literal ℓ_i^j and b its negation bit; then $\Gamma(\bar{i}\bar{j}) = b\bar{k}$ where $\bar{i}, \bar{j}, \bar{k}$ are the respective binary representations of i, j, k . SUCCINCT 3SAT is the problem of deciding, given a circuit Γ , whether the formula φ it encodes is satisfiable.

Let then φ be a formula with 2^m clauses and 2^n variables x_0, \dots, x_{2^n-1} and Γ a boolean circuit encoding this formula. We construct two simple, type-compliant, acyclic processes that are trace equivalent *iff* φ is unsatisfiable. Using pairs $\langle u, v \rangle$ we encode binary trees: a leaf is a non-pair value and, if u and v encode binary trees, $\langle u, v \rangle$ encodes the binary tree whose root has u and v as child nodes. Given a term t , we build a process $P(t)$ behaving as follows:

- (1) $P(t)$ first waits for an input x from the attacker. This term x is expected to be a binary tree of depth n with boolean leaves, modelling a valuation of φ (the i^{th} leaf of x being the valuation of x_i).
- (2) The goal is to make $P(t)$ verify that this valuation satisfies φ ; if the verification succeeds the process outputs t . Given two constants ok and ko , $P(\text{ok})$ and $P(\text{ko})$ will thus be trace equivalent *iff* φ is unsatisfiable.
- (3) However it is not possible to hardcode within a process of polynomial size the verification that the valuation encoded by x satisfies the 2^m clauses of φ . Hence we replicate a process that, given x , verifies one clause at a time. For that for each $0 \leq i \leq 2^m - 1$ we consider the term

$$K_i = \text{senc}(\langle b_0, \langle b_1, \langle \dots b_{m-1} \rangle \dots \rangle \rangle, k)$$

where $b_0 \dots b_{m-1}$ is the binary representation of i and k is a name. Intuitively, the attacker will guide the verification of the 2^m clauses of φ , and whenever the i^{th} clause has been successfully verified, the process reveals the term $\text{senc}(K_i, K_{i-1})$. By convention, K_{-1} is a public constant.

- (4) In particular, the attacker can deduce the term K_{2^m-1} only if she has successfully verified that the initial input x indeed encodes a valuation satisfying all clauses of φ . It therefore suffices to require, before the final output of t , that the attacker inputs K_{2^m-1} . \square

5 COMPARISON WITH OTHER MODELS

In this section we discuss some other notions of indistinguishability and compare them in terms of expressivity and complexity.

5.1 Structure-guided equivalence proofs

We survey two equivalence notions that impose structural constraints to equivalence proofs that make the verification easier.

- The most well-known variant of equivalence properties in security protocols is *diff-equivalence*, different variants of which are proved by the state-of-the-art PROVERIF and TAMARIN. Intuitively, it can be seen as an analogue of trace equivalence where two equivalent traces are also required to follow the exact same execution flow. For example to prove $P_1 \mid \dots \mid P_n$ and $Q_1 \mid \dots \mid Q_n$ equivalent, all actions originated from each subprocess P_i should be matched with actions from Q_i .
- *Equivalence by session* is similar in spirit but impose less restrictions on equivalent traces: rather than sharing the exact same execution flow, they should be organised similarly in terms of parallel sessions. To prove $P_1 \mid \dots \mid P_n$ and $Q_1 \mid \dots \mid Q_n$ equivalent, there should exist a permutation π of $\llbracket 1, n \rrbracket$ such that all actions originated from each P_i should be matched with actions from $Q_{\pi(i)}$. This equivalence has been used in the DEEPSEC tool as a structure-guided heuristic for trace equivalence [CKR19].

Process matchings To formalise this we first define *simplification rules* \rightsquigarrow (Figure 2) that get rid of the deterministic parts of the transition system. They are convergent up to renaming of new names, and we write $P \dot{\rightsquigarrow}$ one arbitrary \rightsquigarrow -normal form of P . A process in \rightsquigarrow -normal form can be uniquely decomposed into

$$P = P_1 \mid \dots \mid P_n = \prod_{i=1}^n P_i \quad (\text{implicit right-associativity})$$

where each P_i starts with an input, an output or a replication.

In order to ensure that two traces verify the structural restrictions of diff-equivalence and equivalence by session, we extend the semantics of the calculus to pair of processes: the rule (PAR) is replaced by a rule pairing parallel subprocesses together, and the communication rules (IN), (OUT), (COMM) can only be triggered when they are applicable to the two components of the pair. Formally this semantics operate on *extended twin processes* $(\mathcal{P}^2, \Phi_0, \Phi_1)$ where \mathcal{P}^2 is a multiset of pairs of processes in \rightsquigarrow -normal form, and Φ_0 and Φ_1 are frames. The semantics of such processes is defined in Figure 3 and assumes that channels are static (they are either constants or names that are never used as parts of output messages) and we use the private semantics (i.e. with no internal communications on public channels). Although one could design a definition making without these two assumptions, they are actively used by the optimisations developed in [CKR19].

The semantics of Figure 3 only handles the bounded fragment, consistently with the presentation of equivalence by session of [CKR19]. However, to avoid being artificially limited in our comparisons, we can naively extend Figure 3 with

$$(\{!P, !Q\} \cup \mathcal{P}^2, \Phi_0, \Phi_1) \xrightarrow{\tau} (\{(!P, !Q), (P \dot{\rightsquigarrow}, Q \dot{\rightsquigarrow})\} \cup \mathcal{P}^2, \Phi_0, \Phi_1)$$

A similar rule can be defined for replication operator if simple pro-

cesses ($!^{ch}$) to bypass the restriction to static channels. This is a natural extension of the semantics, although rather limited too. For example $P \mid !P$ and $!P$ will not be equivalent by session although, intuitively, there exists a natural bijection between all copies of P in $P \mid !P$ and $!P$. We leave open the design of a semantics better adapted to the expected mechanisms of equivalence by session of unbounded processes, and stick to this simplistic model here.

Equivalence by session Two processes P_0 and P_1 are equivalent by session when for all traces of P_i , $i \in \{0, 1\}$, there exists a trace t' of P_{1-i} such that $t \sim t'$ and t and t' are the first and second projections, respectively, of a twin trace of (P_i, P_{1-i}) . In particular equivalence by session refines trace equivalence:

THEOREM 5.1 ([CKR19]). If two processes are equivalent by session then they are also trace equivalent.

The converse is not true in general, consider e.g. the processes $c(x).c(y)$ and $c(x) \mid c(y)$. Besides:

THEOREM 5.2 ([CKR19]). Labelled bisimilarity and equivalence by session are incomparable.

We discuss in Section 6 some assumptions under which trace equivalence and labelled bisimilarity coincide with equivalence by session. We refer to the following problem as SESSSEQ:

INPUT: A theory, two processes

QUESTION: Are the two processes equivalent by session?

Diff equivalence. We formalise diff-equivalence with the same definition as equivalence by session, except that the rule (MATCH) of Figure 3 is restricted to only consider the identity matching:

$$(\{(\prod_{i=1}^n P_i, \prod_{i=1}^n Q_i)\} \cup \mathcal{P}^2, \Phi_0, \Phi_1) \xrightarrow{\tau} (\{(P_i, Q_i)\}_{i=1}^n \cup \mathcal{P}^2, \Phi_0, \Phi_1) \quad (\text{MATCH-ID})$$

Although the original definition of diff-equivalence [BAF08] was stricter by imposing control-flow restrictions on conditionals as well, our formalisation capture a notion similar to the more-relaxed, later-introduced definition of [CB13]. All in all the definition of diff-equivalence, more restrictive than equivalence by session, makes it a sound heuristic all other equivalences:

THEOREM 5.3 ([BAF08, CKR19]). If two processes are diff equivalent then they are also labelled bisimilar, equivalent by session and therefore trace equivalent.

The converse does not hold in general, leading to so-called *false attacks* (non-diff-equivalent processes that are, for example, trace equivalent). They are naturally more frequent than those induced by equivalence by session. Note also that in the extreme case of simple processes, DIFFSEQ and SESSSEQ are essentially the same decision problem, up to a simple associative-commutative preprocessing of parallel operators. We call the following problem DIFFSEQ:

INPUT: A theory, two processes.

QUESTION: Are the two processes diff equivalent?

$$\begin{array}{l}
P \mid 0 \rightsquigarrow P \qquad 0 \mid P \rightsquigarrow P \qquad (P \mid Q) \mid R \rightsquigarrow P \mid (Q \mid R) \qquad \left. \begin{array}{l} P \mid Q \rightsquigarrow P' \mid Q \\ Q \mid P \rightsquigarrow Q \mid P' \end{array} \right\} \text{if } P \rightsquigarrow P' \\
\text{new } k.P \rightsquigarrow P\{k \mapsto k'\} \quad k' \text{ fresh name} \qquad \text{if } u = v \text{ then } P \text{ else } Q \rightsquigarrow \begin{cases} P & \text{if } u =_E v \\ Q & \text{otherwise} \end{cases}
\end{array}$$

Figure 2: Simplification rules for processes

$$\begin{array}{l}
(\{(P, Q)\} \cup \mathcal{P}^2, \Phi_0, \Phi_1) \xrightarrow{\alpha} (\{(P'\xi, Q'\xi)\} \cup \mathcal{P}^2, \Phi'_0, \Phi'_1) \qquad \text{if } (\{P\}, \Phi_0) \xrightarrow{\alpha} (\{P'\}, \Phi'_0), (\{Q\}, \Phi_1) \xrightarrow{\alpha} (\{Q'\}, \Phi'_1) \text{ (IO}^2\text{)} \\
\qquad \text{by rules (IN) or (OUT)} \\
(\{\overline{c_1}\langle u_1 \rangle.P_1, \overline{c_2}\langle u_2 \rangle.P_2, (c_1(x_1).Q_1, c_2(x_2).Q_2)\} \cup \mathcal{P}^2, \Phi_0, \Phi_1) \xrightarrow{\tau} (\{(P_1\xi, P_2\xi), (Q_1\{x_1 \mapsto u_1\}\xi, Q_2\{x_2 \mapsto u_2\}\xi)\} \cup \mathcal{P}^2, \Phi_0, \Phi_1) \qquad \text{if } c_1 \text{ and } c_2 \text{ are private channels} \text{ (COMM}^2\text{)} \\
(\{\prod_{i=1}^n P_i, \prod_{i=1}^n Q_i\} \cup \mathcal{P}^2, \Phi_0, \Phi_1) \xrightarrow{\tau} (\{(P_i, Q_{\pi(i)})\}_{i=1}^n \cup \mathcal{P}^2, \Phi_0, \Phi_1) \qquad \text{if } \pi \text{ is a permutation of } \llbracket 1, n \rrbracket \text{ (MATCH)}
\end{array}$$

Figure 3: Semantics on pairs of processes (in \rightsquigarrow -normal form)

5.2 A tool for decidability: constraint solving

In the bounded fragment it is common to abstract the infinitely-branching transition relation by a finite variant with *symbolic constraints* [Bau07, CCD13, CKR18a], reducing the study of equivalences to various flavours of constraint-solving problems. We detail one of them [Bau07] in this section as it is used in most of the results surveyed in the remaining of the paper.

Constraint systems In a symbolic approach, all recipes are replaced by placeholder variables and constraints are used instead in order to specify how these variables may be instantiated in practice. Typically, see the informal example below:

$$\begin{array}{l}
A = (\{c(x).\text{if } \text{sdec}(x, k) = u \text{ then } \overline{c}\langle x \rangle \text{ else } \overline{c}\langle h(x) \rangle\}, \emptyset) \\
\begin{array}{l} \xrightarrow{Y(X)} (\{\text{if } \text{sdec}(x, k) = u \text{ then } \overline{c}\langle x \rangle \text{ else } \overline{c}\langle h(x) \rangle\}, \emptyset) \\ \xrightarrow{\tau} (\{\overline{c}\langle h(x) \rangle\}, \emptyset) \\ \xrightarrow{\overline{Z}\langle \text{ax} \rangle} (\{\emptyset\}, \Phi) \quad \text{with } \Phi = \{\text{ax} \mapsto h(x)\} \end{array}
\end{array}$$

The three recipes required by the usual semantics are not specified, and three so-called *second-order variables* X, Y, Z are used instead. They may be instantiated by any recipes ξ_X, ξ_Y, ξ_Z that satisfy here the following constraints: ξ_X, ξ_Y, ξ_Z do not use the axiom of Φ , $\xi_Y \Phi =_E c$, $\text{sdec}(\xi_X, k) \Phi \neq_E u$, and $\xi_Z \Phi =_E c$. The set S of these six constraints is usually written

$$S = \{X \vdash^? x, Y \vdash^? y, Z \vdash^? z, \text{sdec}(x, k) \neq^? u, y =^? c, z =^? c\}.$$

The constraint $X \vdash^? x$ is called a *deduction fact* and intuitively indicates that x is deducible by the attacker, using the recipe ξ_X . This recipe may use the first axioms of the frame up to the *arity* of X , written $\text{ar}(X)$. Hence here $\text{ar}(X) = \text{ar}(Y) = \text{ar}(Z) = 0$. The constraints $u =^? v$ (*equations*) and $u \neq^? v$ (*disequations*) express comparisons between terms modulo theory.

Formally, a *constraint system* is a pair $C = (S, \Phi)$ with Φ a frame and S a set of equations, disequations and deduction facts with no

second-order variables appearing twice nor having an arity greater than $|\text{dom}(\Phi)|$. We always assume that they verify the *origination property* which intuitively means that they correspond to actual symbolic traces, i.e. that all free variables appearing in the frame should have been determined by a prior recipe. That is, if

$$\Phi = \{\text{ax}_1 \mapsto t_1, \dots, \text{ax}_n \mapsto t_n\}$$

then the origination property requires that for all $i \in \llbracket 1, n \rrbracket$ and all variables x appearing in t_i , there exists a deduction fact $X \vdash^? x$ in C such that $\text{ar}(X) < i$.

Then a *solution* of a constraint system $C = (S, \Phi)$, substitutes second-order variables by actual recipes that satisfy the equations and disequations of S . Formally a *second-order substitution* is a mapping Σ from second-order variables X to recipes using at most the $\text{ar}(X)$ first axioms of Φ . In particular Σ induces a valuation of the free variables of C , which is the substitution σ such that $X\Sigma = x\sigma$ for all deduction facts $X \vdash^? x$ of S (σ is well-defined and unique under the origination property). We thus say that Σ is a solution of C if $u\sigma \downarrow = v\sigma \downarrow$ for all equations $u =^? v$ of S , and $u\sigma \downarrow \neq v\sigma \downarrow$ for all disequations $u \neq^? v$ of S . In the constructor-destroyer semantics, we additionally require that $\text{msg}(u\sigma)$ and $\text{msg}(v\sigma)$ for an equation to be satisfied, and disequations are satisfied when either $\text{msg}(u\sigma)$ or $\text{msg}(v\sigma)$ does not hold, or $u\sigma \downarrow \neq v\sigma \downarrow$.

Similarly to processes, we say that a constraint system is *positive* when it does not contain disequations.

Constraint solving As we show in the next sections, several equivalence problems are reducible to an analysis of constraint systems, and understanding the complexity of the latter is often key to solve the former. Although its applications are mostly for reachability properties—not surveyed in this paper—we mention the most basic decision problem that we call CSysSAT:

INPUT: a theory, a constraint system.

QUESTION: does the constraint system admit a solution?

This is a generalisation of DEDUCIBILITY since, by definition, a term t is deducible from a frame Φ iff the constraint system

$$(\{X \vdash^? x, x =^? t\}, \Phi) \quad \text{with } ar(X) = |dom(\Phi)|$$

is satisfiable. More generally, the *weak-secrecy* problem (given a process P and a term t , does there exist a trace of P such that t is deducible from its frame?) can be decided in non-deterministic polynomial time with oracle to CSysSAT, intuitively as follows:

- (1) guess non-deterministically one (among the polynomially-many) symbolic execution of P and collect the corresponding constraints into a constraint system $C = (S, \Phi)$
- (2) answer yes if the following constraint system has a solution:

$$(S \cup \{X \vdash^? x, x =^? t\}, \Phi) \quad X, x \text{ fresh, } ar(X) = |dom(\Phi)|.$$

Regarding equivalence properties, the problem is essentially to decide whether two constraint systems admit the same set of solutions, and that their frames are statically equivalent for all of these solutions. In general the decision of trace equivalence involves more complex variants of this decision problem [CCD13, CKR18a], but this simple one is already useful to decide diff-equivalence, as well as other equivalences in some fragments [Bau07, CCD13, CKR19]. We call this problem CSysEq and formalise it as follows:

INPUT: A theory, two constraint systems (S_1, Φ_1) and (S_2, Φ_2) with the same second-order variables and $dom(\Phi_1) = dom(\Phi_2)$.

QUESTION: Do the following two constraint systems have the same set of solutions?

$$(S_1 \cup \{X \vdash^? x, Y \vdash^? y, x =^? y\}, \Phi_1)$$

$$(S_2 \cup \{X \vdash^? x, Y \vdash^? y, x =^? y\}, \Phi_2)$$

with X, Y, x, y fresh variables and $ar(X) = ar(Y) = |dom(\Phi_1)|$.

This problem is called *S-equivalence* in [Bau07]. Note that we retrieve the STATEq problem when S_1 and S_2 are empty.

Complexity We now present some decidability and complexity results for CSysSAT and CSysEq; they will be at the core of the results presented in the next sections. These two problems have been studied in majority in [Bau07] for the decidability of reachability properties and diff equivalence, in the case of *fixed* subterm theories in the positive bounded fragment.

THEOREM 5.4 ([Bau07]). For all fixed subterm convergent theories, CSysSAT (resp. CSysEq) is NP (resp. coNP for positive constraint systems).

As far as we know the complexity of this problem has only been studied for fixed theories. However the result of [Bau07] above can be adapted to parametric theories; inspecting the proof we observe that (1) in the complexity bounds, the dependencies in the theory are polynomial and (2) the proof uses the fact that static equivalence is PTIME for fixed theories (Theorem 3.1) but the arguments still hold if we only assume static equivalence to be coNP. Since it has also been proved in [Bau07] that CSysSAT was NP-hard if the theory includes at least a free binary function symbol, we obtain the more general complexity result:

THEOREM 5.5. CSysSAT (resp. CSysEq) is NP-complete (resp. coNP-complete) for subterm convergent theories and positive constraint systems. In the case of CSysSAT, the NP-completeness also holds without the positivity assumption.

Regarding the complexity lower bounds for fixed theories, similarly to the problems we surveyed in the previous sections, the complexity may vary from one theory to the other. Typically:

THEOREM 5.6. With the empty theory, CSysSAT and CSysEq are LOGSPACE.

Proof. It suffices to prove that CSysEq is LOGSPACE. We let two constraint systems $C_1 = (S_1, \Phi_1)$ and $C_2 = (S_2, \Phi_2)$, where the deduction facts of S_1 and S_2 are, respectively,

$$X_1 \vdash^? x_1, \dots, X_n \vdash^? x_n \quad \text{and} \quad X_1 \vdash^? y_1, \dots, X_n \vdash^? y_n$$

and where $dom(\Phi_1) = dom(\Phi_2) = \{ax_1, \dots, ax_p\}$. In the empty theory, there are finitely-many second-order substitutions Σ for C_1 and C_2 up to bijective renaming of fresh constants (which does not affect whether Σ is a solution of C_1 or C_2). Indeed for all $i \in \llbracket 1, n \rrbracket$, the recipe $X_i \Sigma$ is either

- a constant appearing either in Φ_1, Φ_2 , in an equation of S_1 or S_2 or in some $X_j \Sigma, j < i$
- a fresh constant (i.e. not captured by the previous case)
- an axiom ax_j such that $j < ar(X_i)$.

Given a second-order substitution Σ , we can verify that it is a solution of C_1 and C_2 in LOGSPACE since the constraint systems only contain equations and disequations between constants, names and variables. The problem can thus be solved in LOGSPACE by bruteforce, using three nested loops:

- the first two loops are of size in n and p and are used to enumerate all second-order substitutions Σ up to bijective renaming of fresh constants
- the third loop of size polynomially-bounded by $|C_1| + |C_2|$ verifying that Σ is a solution of C_1 iff it is a solution of C_2 . \square

Since CSysEq is a generalisation of STATEq it can also be interesting to compare their complexity. We already know that they are both LOGSPACE when the empty theory (Theorems 3.2 and 5.6) and that, in the positive fragment, they are both coNP-complete for subterm convergent theories (Theorems 3.4 and 5.5). Regarding fixed theories, STATEq is PTIME (Theorem 3.1) and this is optimal in the sense that the problem is PTIME-hard for all theories containing symmetric encryption (Theorem 3.3). The coNP bound is optimal for CSysEq in the same sense:

THEOREM 5.7. CSysSAT (resp. CSysEq) is NP-hard (resp. coNP-hard) for positive constraint systems if the theory contains at least symmetric encryption.

Proof. It suffices to prove that CSysSAT is NP-hard. By reduction from SAT we let $\varphi = \bigwedge_{i=1}^p C_i$ a SAT formula with variables x_1, \dots, x_n . Given a family of distinct names k_1, \dots, k_n , we first consider the following frame with n free variables

$$\Phi_{val} = \{ax_1 \mapsto \text{senc}(x_1, k_1), \dots, ax_n \mapsto \text{senc}(x_n, k_n)\}.$$

Given a clause C of φ , we let $x_{i_1}, x_{i_2}, x_{i_3}$ its variables, $b_{i_1}, b_{i_2}, b_{i_3}$ its negation bits, and a fresh name k_c . We define a frame Φ_c such that, for all valuations σ of x_1, \dots, x_n , the name k_c is deducible from $\Phi_{val}\sigma \cup \Phi_c$ iff σ satisfies C (i.e. iff there exists $j \in \llbracket 1, 3 \rrbracket$ such that $x_{i_j}\sigma = b_{i_j}$):

$$\Phi_c = \left\{ \begin{array}{l} ax_1^c \mapsto \text{senc}(k_c, \text{senc}(b_{i_1}, k_{i_1})) \\ ax_2^c \mapsto \text{senc}(k_c, \text{senc}(b_{i_2}, k_{i_2})) \\ ax_3^c \mapsto \text{senc}(k_c, \text{senc}(b_{i_3}, k_{i_3})) \end{array} \right\}$$

All in all the following constraint system (S, Φ) is satisfiable iff φ is satisfiable:

$$S = \left\{ \begin{array}{l} X_1 \vdash^? x_1, \dots, X_n \vdash^? x_n, \\ Y_1 \vdash^? y_1, \dots, Y_p \vdash^? y_p, \\ y_1 =^? k_{c_1}, \dots, y_p =^? k_{c_p} \end{array} \right\}$$

$$\Phi = \Phi_{val} \cup \Phi_{c_1} \cup \dots \cup \Phi_{c_p}$$

with $ar(X_1) = \dots = ar(X_n) = 0$, $ar(Y_1) = \dots = ar(Y_p) = |\Phi|$. \square

The complexity of the general problem (that is, with disequations) is open. However it is easily seen less general than trace equivalence and thus inherits its complexity upper bounds.

THEOREM 5.8. CSysEq is reducible to TRACEEq of bounded processes. This reduction is LOGSPACE and preserves the theory.

Proof. Consider a constraint system $C = (S, \Phi)$. We let the notations $\Phi = \{ax_1 \mapsto t_1, \dots, ax_n \mapsto t_n\}$ and $S = D \cup E$ with

$$D = \{X_1 \vdash^? x_1, \dots, X_p \vdash^? x_p\} \quad E = \{u_1 \sim_1 v_1, \dots, u_q \sim_q v_q\}$$

where for all $i \in \llbracket 1, q \rrbracket$, $\sim_i \in \{=^?, \neq^?\}$. Assuming that the second-order variables X_i are sorted by increasing arity, we let

$$1 = i_0 \leq i_1 \leq \dots \leq i_n \leq i_{n+1} = p + 1$$

the sequence of integers such that $ar(X_i) = \ell$ iff $i_\ell \leq i < i_{\ell+1}$. We then let a constant c and define the following process given an other process R :

$$\begin{aligned} P(C, R) &= c(x_{i_0}). \dots c(x_{i_1-1}). \\ &\quad \bar{c}\langle t_1 \rangle. c(x_{i_1}). \dots c(x_{i_2-1}). \\ &\quad \vdots \\ &\quad \bar{c}\langle t_n \rangle. c(x_{i_n}). \dots c(x_{i_{n+1}-1}). \\ &\quad [u_1 \sim_1 v_1] \dots [u_q \sim_q v_q] R \end{aligned}$$

where $[u \sim v]P$ is a shortcut for either “if $u = v$ then P else 0 ” (when \sim is $=^?$) or “if $u = v$ then 0 else P ” (when \sim is $\neq^?$). The process $P(C, R)$ is well-defined (i.e. does not contain variables that are not bound by a prior input) if C verifies the origination property. In $P(C, R)$, the subprocess R can be executed iff x_1, \dots, x_n are instantiated by recipes that define a solution of C . In particular given a constant d and two constraint systems C_0, C_1 verifying the hypotheses of the problem CSysEq, C_0 and C_1 are equivalent iff for all traces t of $P(C_i, \bar{d}\langle d \rangle)$ containing an output on d , $i \in \{0, 1\}$, there exists a trace t' of $P(C_{1-i}, \bar{d}\langle d \rangle)$ such that $t \sim t'$. In particular C_0 and C_1 are equivalent iff

$$P(C_0, \bar{d}\langle d \rangle) + P(C_1, 0) \quad \text{and} \quad P(C_0, 0) + P(C_1, \bar{d}\langle d \rangle)$$

are trace equivalent where, for $k, k' \in \mathcal{N}$ and $e \in \Sigma_0$ fresh:

$$A + B = \bar{e}\langle k \rangle \mid \bar{e}\langle k' \rangle \mid e(x). ([x = k]A \mid [x = k']B) \quad \square$$

COROLLARY 5.9. CSysEq is coNEXP for subterm convergent constructor destructor theories.

5.3 Decidability and complexity

Diff equivalence Although undecidable in general, diff equivalence is decidable in the bounded positive fragment [Bau07]:

THEOREM 5.10 ([Bau07]). In the bounded (resp. bounded positive) fragment, given a non-deterministic algorithm A for non-CSysEq (resp. for non-CSysEq of positive constraint systems), non-DIFFEq is NP, where a call to A is seen as an elementary instruction.

Proof sketch. The decision procedure of [Bau07] for non equivalence consists of (1) guessing a symbolic trace t in one of the processes, (2) consider the unique (if it exists) candidate equivalent trace t' in the other process, and (3) conclude that the processes are not diff-equivalent if the constraint systems corresponding to t and t' are not equivalent. In the case of the positive fragment, an additional argument is required to prove that it is not necessary to consider symbolic traces that produce disequation constraints. \square

In particular when composing this with the different complexity results for CSysEq mentioned in Section 5.2:

COROLLARY 5.11. DIFFEq is coNEXP (resp. coNP) for bounded processes (resp. bounded positive processes) and subterm convergent theories.

The problem is also known coNP-hard even in the positive fragment for a theory containing only a free binary symbol h [Bau07]. However a simple proof justifies that DIFFEq is actually coNP-hard even for the empty theory and, hence, for any fixed theory:

THEOREM 5.12. In the pure pi-calculus, DIFFEq is coNP-complete for positive bounded processes.

Proof. By reduction from SAT let a formula $\varphi = \bigwedge_{i=1}^m C_i$ in CNF and $\vec{x} = x_1, \dots, x_n$ its variables. For each clause C_i , let k_i be a fresh name and define

$$\text{CheckSat}_i(\vec{x}) = [x_{i_1} = b_{i_1}] \bar{c}\langle k_i \rangle \mid \dots \mid [x_{i_p} = b_{i_p}] \bar{c}\langle k_i \rangle$$

where x_{i_1}, \dots, x_{i_p} are the variables of C_i and b_{i_1}, \dots, b_{i_p} their negation bits. That is, at least one output of k_i is reachable in $\text{CheckSat}_i(\vec{x})$ if \vec{x} is a valuation of φ that satisfies C_i . Hence if

$$\text{CheckSat} = c(x_1). \dots c(x_n). (\text{CheckSat}_1(\vec{x}) \mid \dots \mid \text{CheckSat}_m(\vec{x}))$$

$$\text{Final}(t) = c(y_1). [y_1 = k_1] \dots c(y_m). [y_m = k_m] \bar{c}\langle t \rangle$$

then for two distinct constants $0, 1$, $\text{CheckSat} \mid \text{Final}(0)$ and $\text{CheckSat} \mid \text{Final}(1)$ are diff-equivalent iff φ is unsatisfiable. \square

In particular this gives the exact complexity of DIFFEq in the bounded positive fragment. As far as we know the question remains open without the positivity assumption.

COROLLARY 5.13. For subterm convergent theories (fixed or not) and bounded positive processes, DIFFEq is coNP-complete.

Equivalence by session Equivalence by session has been designed as a heuristic to prove trace equivalence by exploiting the structural symmetries that often arise in practical verification. Surprisingly, despite practical improvements by order of magnitudes of the verification time [CKR19], this performance gap is not reflected in the theoretical, worst-case complexity. The same reduction as trace equivalence can indeed be used to prove equivalence by session coNEXP-hard (details in Appendix A).

THEOREM 5.14. There exists a subterm convergent constructor destructor theory for which SESSSEQ is coNEXP-hard for bounded positive processes. Without the positivity requirement, this theory can be limited to symmetric encryption and pairs.

It is discussed in [CKR19] that equivalence by session may also be seen as a standalone security notion in some cases. Intuitively if P, Q are processes operating on a unique channel, proving equivalence by session of $!P$ and $!Q$ means proving trace equivalence of $!^{ch}P$ and $!^{ch}Q$, i.e. the attacker has the capability of distinguishing actions originated from different copies of P or Q . This may be realistic in scenarios where each session of a protocol is dynamically attributed with a port that is observable by the attacker. It appears that this intuition can be formalised, leading to the following result (detailed proof in Appendix C):

THEOREM 5.15. SESSSEQ is reducible to TRACEEQ . This reduction is LOGSPACE and preserves subterm convergence, the constructor-destructor property, boundedness and positiveness.

Proof sketch. Given a theory E and bounded processes P, Q , we construct a theory E' and two bounded processes $\llbracket P \rrbracket, \llbracket Q \rrbracket$ that are trace equivalent w.r.t. E' (in the private semantics) iff P and Q are equivalent by session w.r.t. E .

The theory E' is E extended with pairs and randomised symmetric encryption. Intuitively each parallel subprocess P' of P will start by a fresh renaming of all channels of P' (made public on a parametric channel s) to avoid confusion between the actions of P' and those of other parallel processes. For example a process of the form $P = P_1 \mid \dots \mid P_n$ using a unique public channel c will be encoded as $\llbracket P \rrbracket_s = P'_1 \mid \dots \mid P'_n$ where

$$P'_i = \text{new } c. \text{new } s'. \bar{s}\langle c, s' \rangle. \llbracket P_i \rrbracket_{s'}$$

The encoding of private channels is more involved. We also simulate them using asynchronous public communications, which requires additional modelling tricks to ensure that the output message is not leaked to the adversary and is received by only one parallel private input. For that we use the symmetric encryption of E' to include a mutual-authentication protocol (here the Needham-Schroeder-Lowe protocol) in the encoding of the two communicating processes. \square

In particular TRACEEQ and SESSSEQ have the same complexity for most fragments investigated in this survey. This reduction does not cover the case of the pure pi-calculus, which can however be treated by hand easily, essentially by using the same arguments as for trace equivalence [CKR18a] up to minor changes.

THEOREM 5.16. In the pure pi-calculus, SESSSEQ is Π_2 -complete for bounded processes (resp. bounded positive processes).

6 THE CASE OF DETERMINACY

We now mention the fragment of *determinate* processes, a generalisation of simple processes. In this fragment, most of the studied equivalences coincide and their complexity drops exponentially.

Definition(s) This class has been investigated significantly in the literature [BDH15, CCKK16, CCD13, CKR19] although several variants coexist, as discussed in [BCK20]. For example the results of [BDH15, CKR19] hold for *action-determinate* processes, meaning that they never reach an intermediary state where two inputs (resp. outputs) on the same communication channel are executable in parallel. More formally, given a process P whose channels are all constants, we say that P is action-determinate there exist no traces of either of the following forms:

$$P \xrightarrow{\text{tr}} (\{c(x).P, c(y).P\}) \quad \text{or} \quad P \xrightarrow{\text{tr}} (\{\bar{c}\langle u \rangle.P, \bar{c}\langle v \rangle.P\}).$$

On the other hand a more permissive definition is used in [CCD13] (not detailed in this survey). There also exists a notion that is stricter than all of these, referred as *strong determinacy* [BCK20]. A process is strongly determinate when

- (1) it does not contain private channels,
- (2) it is bounded,
- (3) all its syntactic subprocesses are strongly determinate,
- (4) in case the process is of the form $P \mid Q$ there exist no channels c such that both P and Q contain an input (resp. output) on c .

For example this process is action-determinate but not strongly-determinate:

$$\text{if } a = b \text{ then } c(x) \text{ else } 0 \quad \mid \quad \text{if } a = b \text{ then } 0 \text{ else } c(x).$$

THEOREM 6.1 ([CCD13]). Simple processes are action determinate, and bounded simple processes are strongly determinate.

Effects on the decision of equivalences As mentioned above, the main implication of determinacy is that most equivalences coincide in this fragment:

THEOREM 6.2 ([CCD13, CKR19]). Two labelled bisimilar (resp. equivalent by session) processes are trace equivalent. The converse is true for action-determinate processes.

We recall that, since the model of [CKR19] does not include replication, this theorem is only formally proven in the bounded fragment, regarding equivalence by session. Still, all arguments of [CKR19] carry to our simple extension of equivalence by session to unbounded processes.

Regarding complexity, it is shown in [CCD13] that, for bounded, simple, positive processes, the equivalence problem could be reduced to CSysEQ like Theorem 5.10 for diff-equivalence. Their arguments can be generalised from simple to strongly-determinate processes in a straightforward manner; however it is not clear whether this would also be true for action-determinate processes or processes with else branches. In particular we obtain the same complexity as diff-equivalence for this fragment:

THEOREM 6.3 ([CCD13]). TRACEEQ , BISIM and SESSSEQ are coNP-complete for subterm convergent theories and bounded, strongly determinate, positive processes. The coNP completeness also holds for all fixed subterm convergent theories.

7 SUMMARY AND OPEN PROBLEMS

Table 1 summarises the main results of and highlights remaining open questions (including a few minor results not mentioned in this survey for space reasons, but detailed in the technical report [CKR20b]). Cells for which the complexity results are not tight are colored in grey. For instance, for subterm-convergent constructor-destructor theories and bounded processes, DIFFEq is known coNEXP and coNP -hard, but the precise complexity remains unknown. Consistently with the results of the paper we also include some complexity results with the theory seen as a constant of the problem (denoted as “fixed” in the *theory* columns). The corresponding cells contain bounds applying to *all* theories of the class; e.g. for BISM of bounded processes, with fixed subterm-convergent constructor-destructor theories, the problem is decidable in coNEXP and PSPACE -hard; despite the gap between the two bounds, they are optimal since there exist theories for which the problem is either PSPACE -complete or coNEXP -complete. Therefore this cell is not highlighted in grey. In our opinion the most interesting open questions are:

- Can upper bounds on constructor destructor theories be lifted to more general subterm convergent theories?
- Without the positivity assumption, can we tighten the complexity for diff equivalence, and strongly determinate processes?

This last question might allow to better understand why strongly determinate processes benefit from optimisations that improve verification performance that much. Finally, as witnessed by the contrast between the high complexity of equivalence by session and its practical efficiency, worst-case complexity may not always be an adequate measure.

Acknowledgments The research leading to these result has received funding from the ERC under the EU’s H2020 research and innovation program (grant agreements No 645865-SPOOC), as well as from the French ANR project TECAP (ANR-17-CE39-0004-01). Itsaka Rakotonirina benefits from a Google PhD Fellowship.

REFERENCES

- [ABF17] Martín Abadi, Bruno Blanchet, and Cédric Fournet. The applied pi calculus: Mobile values, new names, and secure communication. *Journal of the ACM (JACM)*, 2017.
- [AC06] Martín Abadi and Véronique Cortier. Deciding knowledge in security protocols under equational theories. *Theoretical Computer Science*, 2006.
- [ACK16] Myrto Arapinis, Véronique Cortier, and Steve Kremer. When are three voters enough for privacy properties? In *European Symposium on Research in Computer Security (ESORICS)*, 2016.
- [ACRR10] Myrto Arapinis, Tom Chothia, Eike Ritter, and Mark Ryan. Analysing unlinkability and anonymity using the applied pi calculus. In *IEEE Computer Security Foundations Symposium (CSF)*, 2010.
- [AF04] Martín Abadi and Cédric Fournet. Private authentication. *Theoretical Computer Science*, 2004.
- [AG99] Martín Abadi and Andrew D Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and computation*, 1999.
- [ANR07] Siva Anantharaman, Paliath Narendran, and Michael Rusinowitch. Intruders with caps. In *International Conference on Rewriting Techniques and Applications*, 2007.
- [BAF08] Bruno Blanchet, Martín Abadi, and Cédric Fournet. Automated verification of selected equivalences for security protocols. *The Journal of Logic and Algebraic Programming*, 2008.
- [Bau07] Mathieu Baudet. *Sécurité des protocoles cryptographiques: aspects logiques et calculatoires*. PhD thesis, 2007.
- [BBK17] Karthikeyan Bhargavan, Bruno Blanchet, and Nadim Kobeissi. Verified models and reference implementations for the TLS 1.3 standard candidate. In *IEEE Symposium on Security and Privacy (S&P)*, 2017.
- [BC14] David A. Basin and Cas Cremers. Know your enemy: Compromising adversaries in protocol analysis. *ACM Transactions on Information and System Security (TISSEC)*, 2014.
- [BCD13] Mathieu Baudet, Véronique Cortier, and Stéphanie Delaune. YAPA: A generic tool for computing intruder knowledge. *ACM Trans. Comput. Log.*, 2013.
- [BCK20] Kushal Babel, Vincent Cheval, and Steve Kremer. On the semantics of communications when verifying equivalence properties. *Journal of Computer Security*, 2020.
- [BDH15] David Baelde, Stéphanie Delaune, and Luca Hirschi. Partial order reduction for security protocols. In *International Conference on Concurrency Theory (CONCUR)*, 2015.
- [BDH⁺18] David A. Basin, Jannik Dreier, Lucca Hirschi, Sasa Radomirovic, Ralf Sasse, and Vincent Stettler. A formal analysis of 5g authentication. In *ACM Conference on Computer and Communications Security (CCS)*, 2018.
- [Bla16] Bruno Blanchet. Modeling and verifying security protocols with the applied pi calculus and proverif. *Foundations and Trends in Privacy and Security*, 2016.
- [CB13] Vincent Cheval and Bruno Blanchet. Proving more observational equivalences with proverif. In *International Conference on Principles of Security and Trust (POST)*, 2013.
- [CBC11] Bruno Conchinha, David A Basin, and Carlos Caleiro. Fast: an efficient decision procedure for deduction and static equivalence. In *International Conference on Rewriting Techniques and Applications (RTA)*, 2011.
- [CC05] Hubert Comon and Véronique Cortier. Tree automata with one memory set constraints and cryptographic protocols. *Theoretical Computer Science*, 2005.
- [CCCK16] Rohit Chadha, Vincent Cheval, Ștefan Ciobăcă, and Steve Kremer. Automated verification of equivalence properties of cryptographic protocols. *ACM Transactions on Computational Logic (TOCL)*, 2016.
- [CCD13] Vincent Cheval, Véronique Cortier, and Stéphanie Delaune. Deciding equivalence-based properties using constraint solving. *Theoretical Computer Science*, 2013.
- [CCD15a] Rémy Chrétien, Véronique Cortier, and Stéphanie Delaune. Decidability of trace equivalence for protocols with nonces. In *IEEE Computer Security Foundations Symposium (CSF)*, 2015.
- [CCD15b] Rémy Chrétien, Véronique Cortier, and Stéphanie Delaune. From security protocols to pushdown automata. *ACM Transactions on Computational Logic (TOCL)*, 2015.
- [CCLD11] Vincent Cheval, Hubert Comon-Lundh, and Stéphanie Delaune. Trace equivalence decision: Negative tests and non-determinism. In *ACM conference on Computer and communications security (CCS)*, 2011.
- [CDD18] Véronique Cortier, Antoine Dallon, and Stéphanie Delaune. Efficiently deciding equivalence for standard primitives and phases. In *European Symposium on Research in Computer Security (ESORICS)*, 2018.
- [CDK09] Ștefan Ciobăcă, Stéphanie Delaune, and Steve Kremer. Computing knowledge in security protocols under convergent equational theories. In *International Conference on Automated Deduction (CADE)*, 2009.

Table 1: Summary of the results. Colored cells indicate configurations with open problems. Naturally, in the case of STATEQ and CSysEQ, the non-applicable hypotheses on processes (e.g. boundedness) should be ignored when reading the table. *All results for diff-equivalence also coincide with the results for trace equivalence, labelled bisimilarity, and equivalence by session for strongly-determinate processes.*

				STATEQ	CSysEQ	DIFFEQ	BISIM	SESEQ	TRACEQ
		<i>theory</i>	<i>process</i>						
subterm convergent		any	bounded		coNP-hard	coNEXP-hard			
			pos.	coNP-complete					
		any (fixed)	bounded		PTIME	coNP-hard	PSPACE-hard	Π_2 -hard	
			pos.	coNP-complete					
	constructor destructor	any	bounded		coNEXP coNP-hard	coNEXP-complete			
			pos.	coNP-complete					
		any (fixed)	bounded		PTIME	coNEXP coNP-hard	coNEXP PSPACE-hard	coNEXP Π_2 -hard	
			pos.	coNP					
	senc, (>)	unbounded	patterned, type compliant, acyclic, simple	PTIME- complete	coNP- complete	PTIME-hard			PRIM REC PTIME-hard
			coNEXP coNP-hard						
		bounded	any				coNEXP PSPACE-hard	coNEXP Π_2 -hard	
		pos.							
empty	bounded	any	LOGSPACE	coNP- complete	PSPACE- complete	Π_2 -complete			
		pos.							

- [CGCG⁺18] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millikan, and Kevin Milner. On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In *ACM Conference on Computer and Communications Security (CCS)*, 2018.
- [CGLM17] Véronique Cortier, Niklas Grimm, Joseph Lallemand, and Matteo Maffei. A type system for privacy properties. In *ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [Che14] Vincent Cheval. Apte: an algorithm for proving trace equivalence. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 2014.
- [CHH⁺17] Cas Cremers, Marko Horvat, Jonathan Hoyland, Sam Scott, and Thyla van der Merwe. A comprehensive symbolic analysis of TLS 1.3. In *ACM Conference on Computer and Communications Security (CCS)*, 2017.
- [CKR18a] Vincent Cheval, Steve Kremer, and Itsaka Rakotonirina. DEEPSEC: Deciding equivalence properties in security protocols theory and practice. In *IEEE Symposium on Security and Privacy (S&P)*, 2018.
- [CKR18b] Vincent Cheval, Steve Kremer, and Itsaka Rakotonirina. The deepsec prover. In *International Conference on Computer Aided Verification (CAV)*, 2018.
- [CKR19] Vincent Cheval, Steve Kremer, and Itsaka Rakotonirina. Exploiting symmetries when proving equivalence properties for security protocols. In *ACM Conference on Computer and Communications Security (CCS)*, 2019.
- [CKR20a] Vincent Cheval, Steve Kremer, and Itsaka Rakotonirina. The hitchhiker’s guide to decidability and complexity of equivalence properties in security protocols. In *Andre Scedrov’s Festschrift (ScedrovFest65)*, 2020.
- [CKR20b] Vincent Cheval, Steve Kremer, and Itsaka Rakotonirina. The hitchhiker’s guide to decidability and complexity of equivalence properties in security protocols (technical report), 2020.
- [Dam97] Mads Dam. On the decidability of process equivalences for the π -calculus. *Theoretical Computer Science*, 1997.
- [DEK82] Danny Dolev, Shimon Even, and Richard M. Karp. On the security of ping-pong protocols. *Information and Control*, 1982.
- [DKR09] Stéphanie Delaune, Steve Kremer, and Mark Ryan. Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security*, 2009.
- [DLM04] Nancy A. Durgin, Patrick Lincoln, and John C. Mitchell. Multiset rewriting and the complexity of bounded security protocols. *Journal of Computer Security*, 2004.
- [DLMS99] Nancy A. Durgin, Patrick Lincoln, John C. Mitchell, and Andre Scedrov. Undecidability of bounded security protocols. In *Proc. Workshop on formal methods in security protocols*, 1999.
- [DY81] D. Dolev and A.C. Yao. On the security of public key protocols. In *Symposium on Foundations of Computer Science (FOCS)*, 1981.
- [FHMS19] Ihor Filimonov, Ross Horne, Sjouke Mauw, and Zach Smith. Breaking unlinkability of the ICAO 9303 standard for e-passports using bisimilarity. In *European Symposium on Research in Computer Security (ESORICS)*, 2019.
- [HS03] Hans Hüttel and Jiri Srba. Recursive ping-pong protocols. *BRICS Report Series*, 2003.
- [Hüt03] Hans Hüttel. Deciding framed bisimilarity. *Electronic Notes in Theoretical Computer Science*, 2003.
- [JK18] Charlie Jacomme and Steve Kremer. An extensive formal analysis of multi-factor authentication protocols. In *IEEE Computer Security Foundations Symposium (CSF)*, 2018.
- [KBB17] Nadim Kobeissi, Karthikeyan Bhargavan, and Bruno Blanchet. Automated verification for secure messaging protocols and their implementations: A symbolic and computational approach. In *IEEE European Symposium on Security and Privacy (EuroS&P)*, 2017.
- [KKNS14] Max I. Kanovich, Tajana Ban Kirigin, Vivek Nigam, and Andre Scedrov. Bounded memory protocols. *Computer Languages, Systems & Structures*, 2014.
- [MPW92] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I. *Inf. Comput.*, 1992.
- [RT03] Michaël Rusinowitch and Mathieu Turuani. Protocol insecurity with a finite number of sessions, composed keys is NP-complete. *Theoretical Computer Science*, 2003.
- [SMCB13] Benedikt Schmidt, Simon Meier, Cas Cremers, and David Basin. The TAMARIN prover for the symbolic analysis of security protocols. In *International Conference on Computer Aided Verification (CAV)*, 2013.

A CO-NEXP HARDNESS OF EQUIVALENCES IN THE BOUNDED FRAGMENT

In this section we prove various statements of coNEXP-hardness of trace equivalence, equivalence by session and labelled bisimilarity (see e.g. Theorems 4.3 and 5.14). As discussed in Section 4.3 we prove the following theorems, that can be seen as extensions of the results of [CKR18a] (that studied the complexity of TRACEEQ and BISIM). First of all we show that the reduction of [CKR18a] can be performed without private channels, at least for trace equivalence and equivalence by session.

THEOREM A.1. There exists a subterm convergent constructor destructor theory such that TRACEEQ and SESSEQ are coNEXP-hard for bounded positive processes whose channels are constants.

Then we also show that this construction can be adapted for the fixed theory of symmetric encryption and pairs, provided we use else branches.

THEOREM A.2. For a theory limited to symmetric encryption and pairs, TRACEEQ and SESSEQ are coNEXP-hard for bounded processes whose channels are constants.

We also obtained the same result for BISIM but this required the use of private channels.

THEOREM A.3. For a theory limited to symmetric encryption and pairs, BISIM is coNEXP-hard for bounded processes.

We prove all these results by reduction from SUCCINCT 3SAT, a NEXP problem that we already described in the body of the paper (see the proof sketch of Theorem 4.6). We thus let φ a formula with 2^m clauses and 2^n variables x_0, \dots, x_{2^n-1} and Γ a boolean circuit encoding this formula.

Proof. We construct two processes P, Q that are equivalent iff φ is unsatisfiable. We observe that circuit computations can be simulated within the positive fragment; for example an OR-gate with two input edges and two output edges would be encoded by

$$d_1(x_1).d_2(x_2). \left(\begin{array}{l} [x_1 = 0] [x_1 = 0] (\bar{e}_1\langle 0 \rangle \mid \bar{e}_2\langle 0 \rangle) \\ | [x_1 = 0] [x_2 = 1] (\bar{e}_1\langle 1 \rangle \mid \bar{e}_2\langle 1 \rangle) \\ | [x_1 = 1] [x_2 = 0] (\bar{e}_1\langle 1 \rangle \mid \bar{e}_2\langle 1 \rangle) \\ | [x_1 = 1] [x_2 = 1] (\bar{e}_1\langle 1 \rangle \mid \bar{e}_2\langle 1 \rangle) \end{array} \right)$$

where d_1, d_2, e_1, e_2 are private channels. Also non-deterministic choice $A + B$, that is a process that is executed either as A or B , can be encoded by:

$$\bar{d}\langle 0 \rangle \mid d(x).A \mid d(x).B$$

where d is a fresh private channel and x a fresh variable. Using these two mechanisms the processes P and Q take the form

$$P = c(x).(CheckTree(x) + CheckSat(x))$$

$$Q = c(x).(CheckTree(x) + CheckSat(x) + Print())$$

The theory covers for a model of binary trees, among others. Intuitively P and Q first receive an input x from the attacker that is expected to be a valuation satisfying φ ; concretely it should be a binary tree of depth n whose leaves are booleans representing the valuation of each variable x_i . Then:

- $Print()$ outputs two messages that will serve as a baseline for comparison with the branches $CheckTree(x)$ and $CheckSat(x)$.
- $CheckTree(x)$ verifies that x is a binary tree of depth n with boolean leaves. To do so it “guesses” an ill-formed branch of x using a sequence of non-deterministic choices; if one is found, $CheckTree(x)$ reaches a state where it outputs messages statically-equivalent to those output by $Print()$.
- Assuming x passes the test of $CheckTree$, $CheckSat(x)$ verifies that the valuation encoded by x satisfies φ . To do so it “guesses” an index $0 \leq i \leq 2^m - 1$ by a sequence of non-deterministic choices, recovers the i^{th} clause of φ by simulating the circuit Γ , and verifies that the valuation encoded by x satisfies this clause. If one non-satisfied clause is found, $CheckSat(x)$ reaches a state where it outputs messages statically-equivalent to those of $Print()$.

All in all, P and Q are equivalent *iff* there exist no valuations x whose encoding by a binary tree passes the test of $CheckSat(x)$, i.e. *iff* φ is unsatisfiable. \square

B CO-NEXP HARDNESS FOR SIMPLE PATTERNED PROCESSES

In this section we prove the coNEXP-hardness of the problem studied in [CCD15a], namely trace equivalence of patterned, simple, acyclic, type-compliant processes with a theory limited to symmetric encryption and pairs (Theorem 4.6). Note that, due to Theorems 6.1 and 6.2, all results for this fragment also apply for labelled bisimilarity, equivalence by session, and therefore diff equivalence.

Proof. We proceed by reduction from Succinct 3SAT (we refer to the sketch of proof of Theorem 4.3 for details about this decision problem). We thus let φ a formula with 2^m clauses and 2^n variables x_0, \dots, x_{2^n-1} and Γ a boolean circuit encoding this formula, and construct two patterned, simple, type-compliant, acyclic processes that are trace equivalent *iff* φ is unsatisfiable.

For succinctness, we use the following notation for tuples of terms $\langle u_1, \dots, u_n \rangle = \langle u_1, \langle u_2, \langle \dots u_n \rangle \dots \rangle \rangle$. Using pairs we can encode binary trees; a leaf is a non-pair value and, if u and v encode binary trees, $\langle u, v \rangle$ encodes the binary tree whose root has u and v as successors. We can then build a process $P(t)$,

where t is a term, that has the following behaviour

- (1) $P(t)$ first waits for an input x from the attacker, expectedly a binary tree of depth n with boolean leaves, modelling a valuation of φ (the i^{th} leaf of x being the valuation of x_i).
- (2) The goal is to make $P(t)$ verify that this valuation satisfies φ ; if the verification succeeds the process will then output m . Given two constants ok and ko , $P(ok)$ and $P(ko)$ will thus be trace equivalent *iff* φ is unsatisfiable.
- (3) However it is not possible to hardcode within a process of polynomial size the verification that the valuation encoded by x satisfies the 2^m clauses of φ . For each $0 \leq i \leq 2^m - 1$ we thus consider a term K_i that depends on i and is unknown to the attacker. For example

$$K_i = \text{senc}(\langle b_0, \dots, b_{m-1} \rangle, k_{\text{reward}})$$

where $b_0 \dots b_{m-1}$ is the binary representation of i . Intuitively the attacker will guide the verification of the 2^m clauses of φ , and the process will reward her by revealing the term $\text{senc}(K_i, K_{i-1})$ when the i^{th} clause has been successfully verified. By convention K_{-1} is a public constant.

- (4) In particular the attacker can deduce the term K_{2^m-1} only if she has successfully verified that the initial input x indeed encodes a valuation of φ that satisfies all clauses of φ . It therefore suffices to require, before the final output of m , that the attacker inputs K_{2^m-1} .

More formally we let a family of names k_a for various labels a , for example $k_{\text{Extract}}, k_{\text{Decr}}, \dots$. The process $P(t)$ then takes the following form; for simplicity we use a single channel c but $P(t)$ can easily be converted to a simple process by using a channel c_R for each parallel subprocess R :

$$P(t) = \text{Extract} \mid \text{Eval} \mid \text{Decr} \mid \text{Init} \mid \text{CheckSat} \mid \text{Final}$$

- Extract performs tree extraction requests: for extracting the component $b \in \{0, 1\}$ of $x = \langle u_0, u_1 \rangle$, the attacker sends a term $\text{senc}(\langle x, b \rangle, k_{\text{Extract}})$ and the process outputs back $\text{senc}(\langle x, b, u_b \rangle, k_{\text{Extract}})$.
- Eval performs circuit-evaluation requests: an input $\vec{b} = b_0, \dots, b_{m+2}$ of Γ is submitted by outputting $\text{senc}(\langle \vec{b} \rangle, k_{\text{Eval}})$ and the process eventually outputs back $\text{senc}(\langle \vec{b}, \vec{c} \rangle, k_{\text{Eval}})$ where $\vec{c} = \Gamma(\vec{b})$.
- Decr performs decrement requests: if the attacker sends a term of the form $\text{senc}(\langle \vec{b} \rangle, k_{\text{Decr}})$ where $\langle \vec{b} \rangle$ is the binary representation of $1 \leq i \leq 2^n - 1$ the process will output back $\text{senc}(\langle \langle \vec{b} \rangle, \langle \vec{b}' \rangle \rangle, k_{\text{Decr}})$ where $\langle \vec{b}' \rangle$ is the binary representation of $i - 1$.
- $\text{Init} = c(x).\bar{c}(\text{senc}(x, k_{\text{CheckSat}}))$ reads a term x from the adversary (supposedly a binary tree encoding a valuation satisfying φ) and locks it under an encryption layer.
- $\text{CheckSat} = !^{\text{ch}} c(\text{senc}(t, k_{\text{CheckSat}})).c(\langle \vec{b} \rangle)$. P receives the term x chosen in Init as well as the binary representation $\langle \vec{b} \rangle$ of an index $0 \leq i \leq 2^m - 1$. The process P then extract the three literals of the i^{th} clause of φ by a sequence of three round-

trip communications with *Eval* with inputs $\langle \vec{b}, 0, 0 \rangle$, $\langle \vec{b}, 0, 1 \rangle$ and $\langle \vec{b}, 1, 0 \rangle$. Say, as a result *Eval* outputs back $\langle b_1 \vec{c}_1 \rangle$, $\langle b_2 \vec{c}_2 \rangle$ and $\langle b_3 \vec{c}_3 \rangle$. The process then attempts to extract from the valuation x the valuation of the variables corresponding to \vec{c}_1 , \vec{c}_2 and \vec{c}_3 which is done by $3n$ round trips with *Extract*. If this is successful, resulting to the three booleans b'_1, b'_2, b'_3 , the i^{th} clause is satisfied if $b_1 = b'_1$, $b_2 = b'_2$ and $b_3 = b'_3$. Under this condition, the process performs a final round trip with *Decr* to compute K_{i-1} , and eventually outputs $\text{senc}(K_i, K_{i-1})$.

- *Final* = $c(K_n). \vec{c}(t)$ verifies that the attacker is able to compute K_n and then outputs t .

This sums up the intuition of the construction. Note that care is needed in order to preserve type compliance and acyclicity. Typically *Decr* has to stratified, that is, $\text{Decr} = !^{ch} \text{Decr}_1 \mid \dots \mid !^{ch} \text{Decr}_m$ where each Decr_i only operates on binary representations of size i . Each Decr_i , $i > 1$, includes a round trip with Decr_{i-1} in case the i^{th} bit is null. A unique replicated process *Decr* that would handle all binary-representation sizes, performing round trips with itself when encountering null bits, would satisfy neither acyclicity nor type-compliance. The same remark applies to *Extract* (one stratum for each tree depth should be defined) and *Eval* (one stratum for each gate of Γ). \square

C REDUCTION OF EQUIVALENCE BY SESSION TO TRACE EQUIVALENCE

In this section we formalise the reduction of the decision of SESS_{EQ} to TRACE_{EQ} (Theorem 5.15).

Proof. Given a theory E and bounded processes P, Q , we construct in polynomial time a theory E' and two bounded processes $\llbracket P \rrbracket, \llbracket Q \rrbracket$ that are trace equivalent w.r.t. E' (in the private semantics) iff P and Q are equivalent by session w.r.t. E .

The theory E' is E extended with pairs and randomised symmetric encryption. Intuitively each parallel subprocess P' of P will start by a fresh renaming of all channels of P' (made public on a parametric channel s) to avoid confusion between the actions of P' and those of other parallel processes. For example a process of the form $P = P_1 \mid \dots \mid P_n$ using a unique public channel c will be encoded as $\llbracket P \rrbracket_s = P'_1 \mid \dots \mid P'_n$ where

$$P'_i = \text{new } c. \text{new } s'. \bar{s} \langle \langle c, s' \rangle \rangle. \llbracket P_i \rrbracket_{s'}$$

The encoding of private channels is more involved. Just as public channels they are dynamically replaced by fresh channels and also revealed to the attacker; as such, internal communications are routed asynchronously through the adversary as well. In particular additional mechanisms are required to get around the following issues:

- Since communications are sent to the attacker, she can get knowledge of their content. In the encoding, messages sent on initially-private channels are therefore encrypted using the fresh encryption function of E' . Note the importance of the encryption to be randomised: otherwise the same message being sent in two internal communications would be reflected as the equality between the ciphertexts observed by

the attacker (which would give her more power compared to synchronous, unobservable communications).

- A naive encoding of internal communications by encrypted public outputs allows replay: the attacker could copy the encrypted message and forward it to several parallel inputs. To avoid this, each private input instructions in the original process is associated with a unique identifier (a constant for example). On the other end, outputs should non-deterministically draw the identifier of an eligible input and initiate a communication protocol (for example the Needham-Schroeder protocol) with it to send the message and ensure that it is not replayed to other identifiers.

The mechanics introduced in the second item prevent private outputs to be received by inputs on different channels, or inputs that are sequentially dependent of the output as in $\vec{c}(u).c(x).P$. \square