

Proving Unlinkability using ProVerif through Desynchronised Bi-Processes

David Baelde
Univ Rennes, CNRS, IRISA, France

Alexandre Debant
Université de Lorraine, CNRS, Inria,
LORIA, F-54000 Nancy, France

Stéphanie Delaune
Univ Rennes, CNRS, IRISA, France

Abstract—Unlinkability is a privacy property of crucial importance for several systems such as mobile phones or RFID chips. Analysing this security property is very complex, and highly error-prone. Therefore, formal verification with machine support is desirable. Unfortunately, existing techniques are not sufficient to directly apply verification tools to automatically prove unlinkability.

In this paper, we overcome this limitation by defining a simple transformation that will exploit some specific features of ProVerif. This transformation, together with some generic axioms, allows the tool to successfully conclude on several case studies. We have implemented our approach, effectively obtaining direct proofs of unlinkability on several protocols that were, until now, out of reach of automatic verification tools.

I. INTRODUCTION

RFID tags are popping up everywhere. They are present in credit cards enabling wireless payments, in wireless keys for opening cars, in e-passport for storing fingerprints and iris scans, in e-ticketing systems for public transport, amusement parks, etc. Citizens do not expect to be traceable just because they carry an e-passport or a credit card. Many applications blatantly ignore such concerns. As a famous and widespread example, contactless credit cards reveal card numbers in clear [3]. More interestingly, subtle tracing attacks may be found even when this security goal is explicitly taken into account. For instance, the BAC protocol for e-passports can be exploited to trace citizens [15]. The attack relies on the possibility for an attacker to trace a passport by observing which error is obtained when replaying a particular message.

According to the ISO/IEC 15408-2 standard, a protocol is said to provide unlinkability when “a user may make multiple uses of [the protocol] without others being able to link these uses together”. This is often defined as the fact that an attacker should not be able to distinguish a scenario in which the same agent (i.e., the user) is involved in many sessions from one that involves different agents in each session. Following this intuition, slightly different definitions have been proposed, e.g. [27], [1], [11], [2]. A comparison between some of them may be found in [12]. In this work, we consider a variant of the definition originally proposed in [1], and revisited in [2] to make it more suitable to analyse RFID protocols manipulating a database. This definition follows the intuition that an attacker

should not be able to distinguish between the real situation where multiple users (e.g. tags) are present and involved in possibly many sessions, and the ideal one in which each user is playing at most one session (in which case it is obviously impossible to link sessions).

Designing protocols which actually achieve the required security goals is well-known to be error-prone. A common good practice is to formally analyse protocols using symbolic techniques and tools, as has been done, e.g., for TLS 1.3 [18], [5]. Aiming at machine support is really relevant since manual proofs are error-prone, tedious and hardly verifiable. Moreover, new protocols are developed quite frequently and need to be verified quickly. These symbolic techniques are mature for reachability properties like confidentiality or authentication, and some tools exist, e.g. Tamarin [25] and ProVerif [8], to ease the verification process. This approach has been extended to deal with privacy properties that are usually expressed through equivalences [9], [4]. Nevertheless, verifying a security property in such a setting, and especially unlinkability as defined above, remains a difficult problem. Actually, ProVerif and Tamarin can only prove a restricted form of equivalence, namely *diff-equivalence*, which (in its standard form) is too limiting to prove unlinkability.

Various avenues have been explored to overcome this limitation and go beyond diff-equivalence, with varying degrees of success. To the best of our knowledge, none of these works is fully satisfactory. Some works rely on a notion of unlinkability expressed using two phases (a learning phase followed by a guessing phase) which is more amenable to automation, but may be weaker than expected. This is the approach followed in [20] where several RFID protocols using the XOR operator have been analysed. However, the security analysis in this work is performed in a restricted setting considering only a fixed number of agents and/or sessions (typically 1 or 2). Another line of works consists of establishing unlinkability using an indirect approach, e.g. [11], [22], [2]. In a nutshell, these works propose conditions that are more easily verified using existing tools, and show that these conditions imply the strongest notion of unlinkability. The latest paper in this line of work, [2], tackles protocols involving a database. It provides an improved notion of unlinkability for these protocols and sufficient conditions that have enabled several first-time unlinkability proofs (as well as the discovery of some attacks). In

the present paper, we shall consider the same class of protocols and the same notion of unlinkability. We will come back to [2] in section V-C to compare its outcomes with the ones resulting from the new approach that we propose. Finally, a natural solution to overcome the limitations of diff-equivalence would be to use *restrictions*, that have been available in Tamarin for a long time. They allow the user to restrict the analysis to a subset of traces, those satisfying the formulas expressed by the restrictions. A rather simple idea would be to apply such a restriction on one side to express the fact that each user can play only once. Unfortunately, restrictions turned out to be incorrectly handled by Tamarin¹ which might introduce a lack of soundness. Recently, the use of restrictions has been formally justified, but for a very specific class (namely Type-0) only [26]. This type of restrictions allows one to enforce protocol phases (used e.g. in the analysis performed in [20]) but cannot capture e.g. equality restrictions needed to model conditional branching. Therefore, the analyses performed in [26] to establish unlinkability revisiting the models provided in [20] have been done considering simplified models (without conditional branching).

Our contributions. The general purpose of our work is to explore how the tool ProVerif can be used to verify unlinkability properties.

We first clarify how the introduction of restrictions in ProVerif [10] can be used to prove meaningful equivalences. This is non-trivial because ProVerif only proves diff-equivalence on bi-processes, and restrictions in that setting formally express properties of bi-executions, while equivalences are more naturally expressed between two independent processes. We identify a class of so-called diff-safe restrictions which, intuitively, have a consistent meaning for the bi-process and its projections. Although this new tool gives us more flexibility to encode security properties, it is still insufficient to model unlinkability.

We then turn to our main contribution. We design transformations of ProVerif models (i.e. bi-processes equipped with axioms and lemmas) which can turn an original model for which ProVerif fails into an equivalent model (in a sense that we make precise later) for which ProVerif succeeds in practice. These transformations are not specific to unlinkability, but we have designed them to tackle the key difficulties faced when trying to automatically verify that property. Our transformations make use of several features which are specific to ProVerif:

- Our first transformation relies on a recent extension of the notion of bi-process. Roughly, a bi-process combines two processes, which are declared equivalent if they can always evolve simultaneously, even for actions that are not observable by the attacker. We propose to dissociate the two processes that form the bi-process when executing some specific unobservable instructions on which the two processes may diverge. This is possible thanks to

an extended kind of bi-process, first considered in [23] and officially introduced in ProVerif in 2020 with version 2.02.

- This transformation alone is not sufficient for the tool to conclude on our case studies because of internal over-approximations. However, adding some generic axioms, as authorised by the latest version of ProVerif presented in [10], improves the accuracy enough to successfully conclude on all our examples.

We have implemented our approach in a tool and applied it on several case studies, obtaining in this way, for the first time, direct proofs of unlinkability for several protocols.

II. MODELLING PROTOCOLS

We recall first how cryptographic messages are modelled as terms in ProVerif. Building on this, we shall see how protocols are modelled as processes, and how protocol indistinguishability is formally expressed through process equivalences.

A. Term algebra

As usual in the symbolic setting, messages are modelled through a term algebra. We assume an infinite set $\mathcal{N}$ of *names* used to represent atomic data such as keys, or nonces, as well as an infinite set Σ_0 of *constants* used to represent atomic data known by the attacker; and two infinite and disjoint sets of *variables* $\mathcal{X}$ and $\mathcal{W}$. Variables in $\mathcal{X}$ are used to refer e.g. to input messages, and variables in $\mathcal{W}$, called *handles*, are used as pointers to messages learned by the attacker. We assume a finite *signature* Σ of *function symbols* with their arity partitioned into two disjoint subsets Σ_c , and Σ_d representing respectively the *constructors*, and *destructors*, and we denote $\Sigma_c^+ = \Sigma_c \uplus \Sigma_0$, and $\Sigma^+ = \Sigma \uplus \Sigma_0$.

We note $\mathcal{T}(\mathcal{F}, D)$ the set of terms built from elements of the set of atomic data D by applying function symbols in the signature $\mathcal{F}$. We refer to elements of $\mathcal{T}(\Sigma^+, \mathcal{N} \cup \mathcal{X})$ as *expressions* and elements of $\mathcal{T}(\Sigma_c^+, \mathcal{N} \cup \mathcal{X})$ as (constructor) terms. We define $\text{vars}(M)$ as the set of variables that occur in an expression M . A term (resp. an expression) is *ground* when it contains no variable. We finally define $\text{names}(M)$ as the set of names occurring in M . We will sometimes use $\text{names}(M)$ and $\text{vars}(D)$ as vectors, meaning that we select an arbitrary ordering over these finite sets.

In order to provide a meaning to constructor symbols, we equip (constructor) terms with an equational theory. We assume a set E of equations over $\mathcal{T}(\Sigma_c, \mathcal{X})$ and we define $=_E$ as the smallest congruence containing E that is closed under substitutions and under bijective renaming. In addition, the semantics of destructor symbols is given by a set R of *ordered rewriting rules* of the form $g(M_1, \dots, M_n) \rightarrow M_0$ with $M_0, M_1, \dots, M_n \in \mathcal{T}(\Sigma_c, \mathcal{X})$. A ground expression D can be rewritten in D' if there is a position p in D , a rewrite rule $g(M_1, \dots, M_n) \rightarrow M_0$ and a substitution θ from variables to ground terms such that $D|_p =_E g(M_1\theta, \dots, M_n\theta)$, and $D' =_E D[M_0\theta]_p$, i.e. D in which the subterm at position p has been replaced by $M_0\theta$. In the case where more than

¹<https://github.com/tamarin-prover/tamarin-prover/issues/324>

one rule may be applied at position p , only the first such rule can be effectively used. Given a ground expression D , it may be possible to rewrite it (in an arbitrary number of steps) into a ground (constructor) term M : in that case, this term is noted $D\Downarrow$, and we say that D *evaluates to* $D\Downarrow$. We write $D\Downarrow$ when no such term exists, and say that the *computation fails*.

For modelling purposes, we split the signature $\Sigma_c \uplus \Sigma_d$ into two parts, Σ_{pub} and Σ_{priv} . An attacker builds his own messages by applying public function symbols to terms he already knows, and which are available to him through variables in $\mathcal{W}$. Formally, a computation done by the attacker is a *recipe*, i.e. a term in $\mathcal{T}(\Sigma_{\text{pub}}^+, \mathcal{W})$ where $\Sigma_{\text{pub}}^+ = \Sigma_{\text{pub}} \uplus \Sigma_0$.

Example 1: We consider $\Sigma_c = \{\langle \rangle, \text{proj}_1, \text{proj}_2, \text{h}\}$, $\Sigma_d = \emptyset$, and we assume that all the function symbols are public. The symbol $\langle \rangle$ of arity 2 is used to model the pairing operator, and we have projection symbols proj_1 and proj_2 (both of arity 1) to access the components of a pair. The symbol h of arity 2 represents a keyed hash function. Given $n_T, k \in \mathcal{N}$, we have that $\langle n_T, \text{h}(n_T, k) \rangle$ is a term.

The standard equational theory E_{pair} modelling the pairing operator contains the two following equations: $\text{proj}_1(\langle x_1, x_2 \rangle) = x_1$, and $\text{proj}_2(\langle x_1, x_2 \rangle) = x_2$.

The informed reader will have noticed that we have chosen here to represent the projections using constructors with an appropriate equational theory. Another choice could have been to use destructors and rewriting rules (instead of equations). Both representations are possible in ProVerif. We made this choice to be consistent with the assumptions of the transformation presented in Section IV-C.

Example 2: To illustrate the notion of destructor symbols and ordered rewriting systems, we consider the signature $\Sigma' = \Sigma'_c \cup \Sigma'_d$ such that $\Sigma'_c = \Sigma_c \cup \{\text{true}, \text{false}\}$ and $\Sigma'_d = \{\text{eq}\}$ together with the rewriting rules $\text{eq}(x, x) \rightarrow \text{true}$, and $\text{eq}(x, y) \rightarrow \text{false}$, in this order (in addition to the equational theory E_{pair} given in Example 1). We assume again that all these function symbols are public.

Let $w \in \mathcal{W}$, and $k_0 \in \Sigma_0$. We have that $R = \text{eq}(\text{h}(\text{proj}_1(w), k_0), \text{proj}_2(w))$ is a recipe, and $D = R\{w \mapsto \langle n_T, \text{h}(n_T, k_0) \rangle\}$, i.e. R in which the occurrences of w are replaced by $\langle n_T, \text{h}(n_T, k_0) \rangle$, is a ground expression such that $D\Downarrow = \text{true}$.

B. Process algebra

We consider two disjoint sets $\mathcal{C}_{\text{pub}}$ and $\mathcal{C}_{\text{priv}}$ of *channel names*. The channels in $\mathcal{C}_{\text{pub}}$ are assumed to be public, whereas those in $\mathcal{C}_{\text{priv}}$ are not. We denote $\mathcal{C} = \mathcal{C}_{\text{pub}} \uplus \mathcal{C}_{\text{priv}}$. We also consider a set Σ_e of *event symbols* and a set Σ_t of *table symbols*. An event symbol evt of arity k may be applied to k terms to form an event $\text{evt}(M_1, \dots, M_k)$. We assume that all table symbols are of arity 1. Events will be used to record particular points in executions, while tables are used to store data across sessions.

Protocols will be modelled as processes using the grammar given in Figure 1. We will not comment on the standard constructs of this grammar (i.e. null process, name and channel

$$\begin{array}{lcl}
 P, Q & := & 0 \\
 & | & \text{new } n; P \\
 & | & \text{out}(c, M); P \\
 & | & \text{new } c_{\text{priv}}; P \\
 & | & \text{in}(c, x); P \\
 & | & \text{let } x = D \text{ in } P \text{ else } Q \\
 & | & \text{event}(e); P \\
 & | & \text{insert } \text{tbl}(M); P \\
 & | & !P \\
 & | & \text{get } \text{tbl}(x) \text{ st. } D \text{ in } P \text{ else } Q \\
 & | & (P \mid Q) \\
 & | & i : P
 \end{array}$$

In this figure, $c \in \mathcal{C}$, $M \in \mathcal{T}(\Sigma_c^+, \mathcal{N} \cup \mathcal{X})$ is a term, $x \in \mathcal{X}$, $n \in \mathcal{N}$, $c_{\text{priv}} \in \mathcal{C}_{\text{priv}}$, $\text{tbl} \in \Sigma_t$, $D \in \mathcal{T}(\Sigma^+, \mathcal{N} \cup \mathcal{X})$ is an expression and e is an event.

Fig. 1. Grammar for processes

restrictions, input, output, parallel composition, replication). The process $\text{insert } \text{tbl}(M); P$ inserts a record in the table tbl then runs P , and $\text{get } \text{tbl}(x) \text{ st. } D \text{ in } P \text{ else } Q$ looks for a record M in the table tbl such that the expression D evaluates to true when x is bound to M . When such a record is found, it runs P in which x has been replaced by M . Otherwise, it runs Q . The construct $\text{let } x = D \text{ in } P \text{ else } Q$ combines a computation with a conditional: it attempts to evaluate D , executes P with x bound to the resulting message upon success, and Q otherwise. The construct $i : P$ models that the process P must be executed in phase i only. Finally, the last construct simply records the event e before executing P .

The constructs in and let bind their variables x in their sub-processes P . The get construct binds x in D and P . The free variables $\text{fv}(P)$ of a process P are defined accordingly. A process P is *closed* when $\text{fv}(P) = \emptyset$.

The operational semantics of processes is given by a labelled transition system over configurations (denoted by K) representing the current state of the process.

Definition 1: A *configuration* is a tuple $(\mathcal{P}; \Phi; \mathcal{S}; i)$ with $i \in \mathbb{N}$ and such that:

- $\mathcal{P}$ is a multiset of closed processes;
- $\Phi = \{w_1 \mapsto M_1, \dots, w_n \mapsto M_n\}$ is a *frame* representing the knowledge of the adversary, i.e. a substitution where $w_1, \dots, w_n$ are handles and $M_1, \dots, M_n$ are ground terms;
- $\mathcal{S}$ is a set of elements of the form $\text{tbl}(M)$ with $\text{tbl} \in \Sigma_t$, and M a ground term. This set $\mathcal{S}$ represents the content of all the tables.

We may write P for the configuration $(\{0 : P\}; \emptyset; \emptyset; 0)$, where we explicitly indicate that the process is intended to execute in phase 0.

The labelled transitions are defined using usual rules (see e.g. [17], [10]). Some rules are recalled in Figure 2. An adversary can send any message of their knowledge on a public channel through the IN rule, while communications are

IN	$(\{i : \text{in}(c, x); P\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\text{in}(c, R)} (\{i : P\{x \mapsto M\}\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i)$ if $c \in \mathcal{C}_{\text{pub}}$ and there exists a recipe R such that $R\Phi \Downarrow =_E M$
IO	$(\{i : \text{in}(c, x); P\} \uplus \{i : \text{out}(c, M); Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau} (\{i : P\{x \mapsto M\}\} \uplus \{i : Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i)$ if $c \in \mathcal{C}_{\text{priv}}$
INSERT	$(\{i : \text{insert } \text{tbl}(M); P\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau} (\{i : P\} \uplus \mathcal{P}; \Phi; \mathcal{S} \cup \{\text{tbl}(M)\}; i)$
GET-THEN	$(\{i : \text{get } \text{tbl}(x) \text{ st. } D \text{ in } P \text{ else } Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau} (\{i : P\{x \mapsto M\}\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i)$ if there exists M such that $\text{tbl}(M) \in \mathcal{S}$, and $D\{x \mapsto M\}$ evaluates to true
GET-ELSE	$(\{i : \text{get } \text{tbl}(x) \text{ st. } D \text{ in } P \text{ else } Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau} (\{i : Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i)$ if for any M such that $\text{tbl}(M) \in \mathcal{S}$, $D\{x \mapsto M\}$ does not evaluate to true
MOVE	$(\mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\text{phase}(i+1)} (\mathcal{P}; \Phi; \mathcal{S}; i+1)$
PHASE	$(\{i : i' : P\} \cup \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau} (\{i' : P\} \cup \mathcal{P}; \Phi; \mathcal{S}; i)$

Fig. 2. Semantics for processes (selected rules)

synchronous on private channels (rule IO). Regarding tables, the semantics is given by the rules INSERT, GET-THEN, and GET-ELSE. We also have rules for the other constructs, and some rules to deal with phases, e.g. MOVE and PHASE. An *execution trace* is a finite sequence:

$$K_0 \xrightarrow{\ell_1} K_1 \xrightarrow{\ell_2} \dots \xrightarrow{\ell_n} K_n.$$

The associated sequence of observables of such a trace is the sequence $\ell_1, \dots, \ell_n$ from which τ actions have been removed. Given a closed process P , we denote $\text{traces}(P)$ the set of execution traces from P . Given an integer j such that $0 \leq j \leq n$, we denote $T[j]$ the configuration K_j , i.e. the configuration at step j in trace T .

Example 3: For illustration purposes, we consider the Basic Hash protocol as described in [11]. Each tag stores a secret key that is never updated, and the readers have access to a database containing all these keys. We have that:

$$T \rightarrow R : \langle n_T, h(n_T, k) \rangle$$

where n_T is a fresh nonce and k is the secret key. When receiving a message, the reader checks that it is a pair whose second element is a hash of the first element of the pair with one of the keys from the database. The protocol can be modelled in our syntax as follows:

$$P_{\text{BH}} = (0 : ! \text{new } k; \text{insert } \text{keys}(k); 1 : ! P_{\text{T}}) \mid (1 : ! P_{\text{R}})$$

where $k, n_T \in \mathcal{N}$, $y, z \in \mathcal{X}$, $\text{ok}, \text{error} \in \Sigma_0$, $\text{keys} \in \Sigma_t$, and $c_R, c_T \in \mathcal{C}_{\text{pub}}$. Moreover, we have that:

- $P_{\text{R}} = \text{in}(c_R, x);$
 $\text{get } \text{keys}(y) \text{ st. } D \text{ in out}(c_R, \text{ok}) \text{ else out}(c_R, \text{error});$
- $D = \text{eq}(\text{proj}_2(x), h(\text{proj}_1(x), y));$ and
- $P_{\text{T}} = \text{new } n_T; \text{out}(c_T, \langle n_T, h(n_T, k) \rangle).$

C. Trace equivalence

We now define the notion of *trace equivalence* on which our definition of unlinkability is based. Intuitively, two processes

are trace equivalent if for each trace of one process, there is an indistinguishable trace of the other process. To define this formally, we first introduce *static equivalence* between frames. Intuitively, an attacker can distinguish two frames Φ and Φ' if there exists a test that fails in Φ and succeeds in Φ' (or the contrary).

Definition 2: Two frames Φ and Φ' are in *static equivalence*, written $\Phi \sim_s \Phi'$, when $\text{dom}(\Phi) = \text{dom}(\Phi')$ and for any recipes R_1 and R_2 , we have that:

$$R_1\Phi \Downarrow =_E R_2\Phi \Downarrow, \text{ if, and only if, } R_1\Phi' \Downarrow =_E R_2\Phi' \Downarrow.$$

Example 4: Continuing with the Basic Hash protocol, we consider the two following frames $(n_T, n'_T, k, k' \in \mathcal{N})$:

- $\Phi_{\text{diff}} = \{w_1 \mapsto \langle n_T, h(n_T, k) \rangle; w_2 \mapsto \langle n'_T, h(n'_T, k') \rangle\};$
and
- $\Phi_{\text{same}} = \{w_1 \mapsto \langle n_T, h(n_T, k) \rangle; w_2 \mapsto \langle n'_T, h(n'_T, k) \rangle\}.$

We have that $\Phi_{\text{diff}} \sim_s \Phi_{\text{same}}$. Intuitively, an attacker cannot distinguish two outputs from the same tag from two outputs performed by different tags.

We then extend the notion of equivalence to traces as follows: given two execution traces T and T' leading respectively to configurations $K = (\mathcal{P}; \Phi; \mathcal{S}; i)$ and $K' = (\mathcal{P}'; \Phi'; \mathcal{S}'; i')$, we write $T \cong T'$ when $i = i'$, T and T' have the same sequence of observables, and $\Phi \sim_s \Phi'$.

Finally, we define the notion of *trace equivalence* of two closed processes by considering all their execution traces.

Definition 3: Let P and P' be two closed processes. We say that P and P' are in *trace equivalence*, written $P \approx_t P'$, when, for any execution trace $T \in \text{traces}(P)$, there exists $T' \in \text{traces}(P')$ such that $T \cong T'$, and conversely.

We rely on this notion of trace equivalence to model unlinkability following the definition originally proposed in [1] and used in [21], [24] relying on a notion of equivalence stronger than the notion of trace equivalence. Actually, the notion of unlinkability we consider here is exactly the one studied in [2]. The purpose of this work is not to provide a discussion on how to model unlinkability but to show how to overcome the

limitations of the existing tools such ProVerif and Tamarin, and to show how to automate proofs of unlinkability in ProVerif relying on some recent advanced features added into the tool. Therefore, we simply illustrate the definition we will use here on the Basic Hash example.

Example 5: Going back to our running example, unlinkability will be modelled as the following equivalence:

$$P_{\text{BH}} \stackrel{?}{\approx}_t (0 : ! \text{new } k; \text{insert } \text{keys}(k); 1 : P_{\text{T}}) \mid (1 : ! P_{\text{R}})$$

This equivalence states that an attacker cannot distinguish the situation where many tags play multiple sessions from the situation where each tag can play only once. Note that the replication $!$ in front of process P_{T} has been removed on the right hand side of the equivalence.

Tools like Tamarin or ProVerif cannot currently prove this equivalence, as observed in [2]. We explain how ProVerif works and why it fails on our example in Section III. We then describe our approach to solve this issue in Section IV.

III. PROVING EQUIVALENCES USING PROVERIF

ProVerif can prove an equivalence between two processes P and Q that differ only by the terms and expressions they contain. In Section III-A, we explain this notion of bi-processes, and the notion of diff-equivalence that ProVerif considers. Then, we describe some features that have been recently introduced in ProVerif [10] and on which our approach will rely.

A. Bi-processes and diff-equivalence

Two processes that only differ by terms and expressions can be given by a single process in which the special function symbol $\text{diff}[\cdot, \cdot]$ is used. Intuitively, the first argument corresponds to the term used in the first process, whereas the second one is the term used in the second process. More formally, the grammar we consider for bi-processes is the same as the one introduced previously with the addition of $\text{diff}[M, M']$ for terms and the systematic use of $\text{diff}[x, y]$ for variable bindings in let, get, and input commands. For instance an input bi-process is of the form $\text{in}(c, \text{diff}[x, y]); P$. The latter addition means that we do not consider ProVerif's standard bi-processes but their extension allowed by the `allowDiffPatterns` option.

When P is a bi-process we define $\text{fst}(P)$ as the process obtained by replacing any subterm $\text{diff}[M, M']$ by M in P . We define similarly $\text{fst}(M)$ when M is a bi-term, and define symmetrically $\text{snd}(P)$ and $\text{snd}(M)$. ProVerif's standard bi-processes, without diff operators on variable bindings, can equivalently be characterised in our setting as *separated* bi-processes. To formally define this notion, we assume given a partition $\mathcal{X} = \mathcal{X}^{\text{L}} \uplus \mathcal{X}^{\text{R}}$ for variables.

Definition 4: A bi-process B is *separated* when its variable bindings are of the form $\text{diff}[x^{\text{L}}, x^{\text{R}}]$ with $x^{\text{L}} \in \mathcal{X}^{\text{L}}$ and $x^{\text{R}} \in \mathcal{X}^{\text{R}}$, and all (bi)terms M occurring in B are such that $\text{vars}(\text{fst}(M)) \subseteq \mathcal{X}^{\text{L}}$ and $\text{vars}(\text{snd}(M)) \subseteq \mathcal{X}^{\text{R}}$.

When a closed bi-process P is separated, $\text{fst}(P)$ and $\text{snd}(P)$ are closed processes too.

Example 6: Let $B = \text{in}(c, \text{diff}[x^{\text{L}}, x^{\text{R}}]); \text{out}(c, \langle x^{\text{L}}, x^{\text{R}} \rangle)$. We have that B is a closed bi-process since the variables x^{L} and x^{R} are bound by the input construct, but it is not separated. Indeed, we have that $M = \langle x^{\text{L}}, x^{\text{R}} \rangle$ occurs in B , and $\text{vars}(\text{fst}(M)) = \text{vars}(\text{snd}(M)) = \{x^{\text{L}}, x^{\text{R}}\}$. Thus, we have neither $\text{vars}(\text{fst}(M)) \subseteq \mathcal{X}^{\text{L}}$ nor $\text{vars}(\text{snd}(M)) \subseteq \mathcal{X}^{\text{R}}$. Moreover, $\text{fst}(B) = \text{in}(c, x^{\text{L}}); \text{out}(c, \langle x^{\text{L}}, x^{\text{R}} \rangle)$ is not closed. Note that $\text{snd}(B)$ is not closed as well.

The semantics of bi-processes is defined by a labelled transition system over bi-configurations of the form $(\mathcal{P}; \Phi; \mathcal{S}; i)$ where:

- $\mathcal{P}$ is a multiset of closed bi-processes,
- Φ is a substitution mapping variables to bi-terms called a *bi-frame*,
- $\mathcal{S}$ is a *bi-store*, i.e. a set of items of the form $\text{tbl}(M)$ where tbl is a table identifier and M is a bi-term, and
- i is an integer.

This transition system over bi-configurations is actually an adaptation of the transition system for regular configurations. Some rules are given in Figure 3. When Φ is a bi-frame, we write Φ_{fst} for the frame of domain $\text{dom}(\Phi)$ such that $\Phi_{\text{fst}}(w) = \text{fst}(\Phi(w))$. For a bi-store $\mathcal{S}$, we define $\mathcal{S}_{\text{fst}} = \{\text{tbl}(\text{fst}(M)) \mid \text{tbl}(M) \in \mathcal{S}\}$. We define similarly Φ_{snd} and $\mathcal{S}_{\text{snd}}$. Note that, in this semantics, both sides of the bi-process have to progress in the same way. For instance, when executing a let instruction, the left and right processes have to progress in the same branch. When executing a get instruction, the left and right processes have to execute the same branch when they access to the same element.

For separated bi-processes, $(\mathcal{P}; \Phi; \mathcal{S}; i) \rightarrow_b (\mathcal{P}'; \Phi'; \mathcal{S}'; i)$ implies $(\text{fst}(\mathcal{P}); \Phi_{\text{fst}}; \mathcal{S}_{\text{fst}}; i) \rightarrow (\text{fst}(\mathcal{P}'); \Phi'_{\text{fst}}; \mathcal{S}'_{\text{fst}}; i)$ and similarly for the second projection. This does not hold in general and, as illustrated by Example 6, it may not even make sense as the projections are not necessarily closed processes.

An execution trace of a bi-process is called a *bi-execution trace*. Given a bi-execution

$$T = K_0 \xrightarrow{\ell_1}_b K_1 \xrightarrow{\ell_2}_b \dots \xrightarrow{\ell_n}_b K_n$$

and an integer $0 \leq j \leq n$, we denote by $T[j]$ the bi-configuration at step j . As for processes, we may simply write B for the bi-configuration $(\{0 : B\}; \emptyset; \emptyset; 0)$.

The notion of convergence defined next captures the fact that the left and right processes always agree along a trace, in the sense that one of the above rules is always applicable unless both sides of the bi-process are blocked.

Definition 5: A bi-execution trace T *converges*, denoted $T \Downarrow$, when for all steps j , $T[j] = K = (\mathcal{P}; \Phi; \mathcal{S}; i)$ implies:

- 1) if $(i : \text{get } \text{tbl}(\text{diff}[x^{\text{L}}, x^{\text{R}}]) \text{ st. } D \text{ in } P \text{ else } Q) \in \mathcal{P}$ then for all $\text{tbl}(M) \in \mathcal{S}$, we have that:

$$\text{fst}(D)\{x^{\text{L}} \mapsto \text{fst}(M), x^{\text{R}} \mapsto \text{snd}(M)\} \Downarrow_{\text{E}} \text{true}$$

if, and only if,

$$\begin{array}{ll}
\text{IN}_b & (\{i : \text{in}(c, \text{diff}[x^L, x^R]); P\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \\
& \xrightarrow{\text{in}(c, R)} (\{i : P\{x^L \mapsto M, x^R \mapsto M'\}\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \\
& \text{if } c \in \mathcal{C}_{\text{pub}} \text{ and } R\Phi_{\text{fst}} \Downarrow =_{\text{E}} M \text{ and } R\Phi_{\text{snd}} \Downarrow =_{\text{E}} M' \\
\\
\text{IO}_b & (\{i : \text{in}(c, \text{diff}[x^L, x^R]); P\} \uplus \{i : \text{out}(c, M); Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \\
& \xrightarrow{\tau} (\{i : P\{x^L \mapsto \text{fst}(M), x^R \mapsto \text{snd}(M)\}\} \uplus \{i : Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \text{ if } c \in \mathcal{C}_{\text{priv}} \\
\\
\text{LET-THEN}_b & (\{i : \text{let diff}[x^L, x^R] = D \text{ in } P \text{ else } Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau_b} \\
& (\{i : P\{x^L \mapsto \text{fst}(D) \Downarrow, x^R \mapsto \text{snd}(D) \Downarrow\}\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \\
\\
\text{LET-ELSE}_b & (\{i : \text{let diff}[x^L, x^R] = D \text{ in } P \text{ else } Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau_b} (\{i : Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \\
& \text{if } \text{fst}(D) \Downarrow \text{ and } \text{snd}(D) \Downarrow \\
\\
\text{GET-THEN}_b & (\{i : \text{get tbl}(\text{diff}[x^L, x^R]) \text{ st. } D \text{ in } P \text{ else } Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau_b} \\
& (\{i : P\{x^L \mapsto \text{fst}(M), x^R \mapsto \text{snd}(M)\}\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \\
& \text{if there exists } \text{tbl}(M) \in \mathcal{S} \text{ such that } \text{fst}(D)\{x^L \mapsto \text{fst}(M), x^R \mapsto \text{snd}(M)\} \Downarrow =_{\text{E}} \text{true} \\
& \text{and } \text{snd}(D)\{x^L \mapsto \text{fst}(M), x^R \mapsto \text{snd}(M)\} \Downarrow =_{\text{E}} \text{true} \\
\\
\text{GET-ELSE}_b & (\{i : \text{get tbl}(\text{diff}[x^L, x^R]) \text{ st. } D \text{ in } P \text{ else } Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \xrightarrow{\tau_b} (\{i : Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}; i) \\
& \text{if for all } \text{tbl}(M) \in \mathcal{S}, \text{fst}(D)\{x^L \mapsto \text{fst}(M), x^R \mapsto \text{snd}(M)\} \text{ does not evaluate to true} \\
& \text{and } \text{snd}(D)\{x^L \mapsto \text{fst}(M), x^R \mapsto \text{snd}(M)\} \text{ does not evaluate to true}
\end{array}$$

Fig. 3. Semantics for bi-processes

- $\text{snd}(D)\{x^L \mapsto \text{fst}(M), x^R \mapsto \text{snd}(M)\} \Downarrow =_{\text{E}} \text{true};$
- 2) $(i : \text{let diff}[x^L, x^R] = D \text{ in } P \text{ else } Q) \in \mathcal{P}$ then $\text{fst}(D)$ evaluates if, and only if, $\text{snd}(D)$ evaluates;
 - 3) $\text{fst}(\Phi) \sim_s \text{snd}(\Phi)$.

We naturally extend this notion to sets of bi-traces: $\mathcal{T} \Downarrow \uparrow$ if, and only if, $T \Downarrow \uparrow$ for all $T \in \mathcal{T}$. We finally say that diff-equivalence holds for a bi-process B when $\text{traces}(B) \Downarrow \uparrow$.

In [9], Blanchet *et al.* proved an important result for separated bi-processes: if $\text{traces}(B)$ converges, i.e. $\text{traces}(B) \Downarrow \uparrow$, then $\text{fst}(B)$ is observationally equivalent to $\text{snd}(B)$. In short, and because observational equivalence implies trace equivalence [13], we have:

if B is separated and $\text{traces}(B) \Downarrow \uparrow$ then $\text{fst}(B) \approx_t \text{snd}(B)$

Example 7: Going back to our running example, we form the bi-process $B_{\text{BH}} =$

$$(0 : ! \text{new } k^L; ! \text{new } k^R; \text{insert keys}(\text{diff}[k^L, k^R]); 1 : B_{\text{T}}) \\
| (1 : ! B_{\text{R}})$$

where $B_{\text{T}} = P_{\text{T}}\{k \mapsto \text{diff}[k^L, k^R]\}$, and B_{R} is obtained from P_{R} by changing x (resp. y) into $\text{diff}[x^L, x^R]$ (resp. $\text{diff}[y^L, y^R]$). This is a separated bi-process.

We can recognise the two processes compared in Example 5 as the projections of our bi-process. Indeed, when considering the first projection $\text{fst}(B_{\text{BH}})$, and after having removed the name k^R that is not used in the projected process, we get

$$(0 : ! \text{new } k^L; ! \text{insert keys}(k^L); 1 : \text{fst}(B_{\text{T}})) \\
| (1 : ! \text{fst}(B_{\text{R}})).$$

which corresponds to P_{BH} in Example 3 (note that $\text{fst}(B_{\text{T}}) = P_{\text{T}}\{k \mapsto k^L\}$ and $\text{fst}(B_{\text{R}}) = P_{\text{R}}\{k \mapsto k^R\}$). The unique difference is that the replication command is executed before

the table insertion instead of before the process P_{T} . This has no significant impact; there will just be many instances of the same key k^L in the table. Similarly, when considering the second projection $\text{snd}(B_{\text{BH}})$, and after having removed the unused name k^L , we get

$$(0 : ! \text{new } k^R; \text{insert keys}(k^R); 1 : \text{snd}(B_{\text{T}})) \\
| (1 : ! \text{snd}(B_{\text{R}}))$$

which corresponds to the right-hand side process in Example 5. Semantically, the two consecutive replication operators are indeed equivalent to a single one.

In conclusion we have that $\text{fst}(B_{\text{BH}})$ represents the situation with multiple tags playing multiple sessions, while $\text{snd}(B_{\text{BH}})$ represents tags that play only once.

Unfortunately, the set $\text{traces}(B_{\text{BH}})$ of bi-execution traces is not convergent. Indeed, as explained in [2], starting from B_{BH} , we can reach a bi-configuration $(\mathcal{P}; \Phi; \mathcal{S}; i)$ such that $\mathcal{P}$ contains the bi-process

$$1 : \text{get keys}(\text{diff}[y^L, y^R]) \text{ st. } D\sigma \text{ in } \dots \text{ else } \dots$$

where:

- $\sigma = \{x^L \mapsto \text{fst}(M), x^R \mapsto \text{snd}(M)\}$,
- $M = \langle n_T, h(n_T, \text{diff}[k^L, k_1^R]) \rangle$, and
- $\text{keys}(\text{diff}[k^L, k_1^R]), \text{keys}(\text{diff}[k^L, k_2^R]) \in \mathcal{S}$.

Considering the element $\text{keys}(\text{diff}[k^L, k_2^R]) \in \mathcal{S}$, we can see that item 1 of Definition 5 is not satisfied.

Although Tamarin relies on a different approach to establish equivalences, a similar encoding adapted to the syntax of Tamarin will lead to the exact same issue (see file BasicHash-Tamarin.spthy in our repository).

B. Analysing bi-processes with restrictions, lemmas and axioms

An extension of ProVerif with restrictions, lemmas, and axioms has recently been introduced [10]. We will see that these features (especially axioms) are actually useful to establish unlinkability.

We consider *formulas* and *correspondence queries* given by the grammar below, where $e, e_1, \dots, e_n$ are events and M, N are (constructor) terms without `diff` operators:

$$\begin{aligned} \psi, \psi' &:= \text{true} \mid \text{false} \mid \text{event}(e) \mid M = N \mid M \neq N \\ &\quad \mid \psi \wedge \psi' \mid \psi \vee \psi' \\ \rho &:= \text{event}(e_1) \wedge \dots \wedge \text{event}(e_n) \Rightarrow \psi \end{aligned}$$

Intuitively, a trace T satisfies the correspondence query ρ (noted $T \vdash \rho$) if whenever T contains instances of `event`(e_i) at τ_i for each $1 \leq i \leq n$, then T also satisfies ψ (for some instantiation of the variables that occur in ψ but not in the e_i) where each event in ψ is found in T at some step τ that occurs before some τ_i . A closed process P satisfies ρ when $T \vdash \rho$ for any $T \in \text{traces}(P)$.

This notion of correspondence queries is extended straightforwardly to bi-execution traces by considering that events contain bi-terms. The formal definition of satisfaction is actually the same as the one given above. We want to emphasise that this formal definition is different from the more natural one that would consist in checking satisfaction of each projection of the bi-restriction by the corresponding projection of the bi-trace. We first illustrate this on an example, and we will come back to this point in Section III-C.

Example 8: Consider the following bi-process B :

```
! new  $k^L$ ; new  $k^R$ ;
! new  $sid$ ; event( $e(sid, \text{diff}[k^L, k^R])$ ); out( $c, \text{diff}[k^L, k^R]$ ).
```

Let ρ_2 be the following correspondence query:

$$\begin{aligned} &\text{event}(e(\text{diff}[x_1, x_2], \text{diff}[y_1, y_2])) \\ &\wedge \text{event}(e(\text{diff}[x'_1, x'_2], \text{diff}[y'_1, y'_2])) \Rightarrow x_2 = x'_2 \end{aligned}$$

We have that $\text{traces}(B) \not\vdash \rho_2$. Indeed, consider

$$T = \left\{ \begin{array}{l} \text{event}(e(sid_1, \text{diff}[k^L, k^R])); \\ \text{event}(e(sid_2, \text{diff}[k^L, k^R])); \end{array} \right.$$

We have that $T \in \text{traces}(B)$ and $T \not\vdash \rho_2$ as $sid_1 \neq sid_2$.

Correspondence queries may be used in several ways in the verification of a bi-process. We call *restriction* a correspondence query that is going to be used to restrict the set of bi-execution traces that we want to consider. We call *axiom* a correspondence query that holds for all bi-execution traces, and we call it a *lemma* when this fact should be verified by ProVerif instead of just being assumed. The next theorem sums up what it means to verify a bi-process in presence of axioms, lemmas and restrictions using ProVerif's semi-decision procedure `prove'` for equivalence queries — axioms and lemmas are similarly handled in this theorem, but in practice axioms would be trusted while lemmas would be verified.

Given a closed bi-process B and a set of restrictions $\mathcal{R}$, we note $\text{traces}^{\mathcal{R}}(B) = \text{traces}(B) \cap \{T \mid T \vdash \rho, \forall \rho \in \mathcal{R}\}$. Moreover, we note $\text{fst}(\mathcal{R}) = \{\text{fst}(\rho) \mid \rho \in \mathcal{R}\}$ and $\text{snd}(\mathcal{R}) = \{\text{snd}(\rho) \mid \rho \in \mathcal{R}\}$.

Example 9: Going back to Example 8, we have that:

$$\begin{aligned} \text{fst}(\rho_2) &= \text{event}(e(x_1, y_1)) \wedge \text{event}(e(x'_1, y'_1)) \Rightarrow x_2 = x'_2 \\ \text{snd}(\rho_2) &= \text{event}(e(x_2, y_2)) \wedge \text{event}(e(x'_2, y'_2)) \Rightarrow x_2 = x'_2 \end{aligned}$$

It is important to note that we have $\text{snd}(T) \not\vdash \text{snd}(\rho_2)$ as expected but $\text{fst}(T) \vdash \text{fst}(\rho_2)$. Indeed, according to ProVerif semantics, variables that appear on the conclusion of a query (or lemma and restriction) and that are not introduced in its premises, are existentially quantified. This means that $\text{fst}(\rho_2)$ is satisfied by any trace. We wanted to point out this (maybe unexpected) semantics feature of ProVerif to ease the understanding of some examples.

Theorem 1 ([10]): Let B be a closed bi-process, $\mathcal{A}x, \mathcal{L}, \mathcal{R}$ be respectively sets of axioms, lemmas, and restrictions. If

- for all $\varrho \in \mathcal{A}x \cup \mathcal{L}$, we have that $\text{traces}^{\mathcal{R}}(B) \vdash \varrho$; and
- `prove'(B, Ax, L, R)` returns true

then $\text{traces}^{\mathcal{R}}(B) \downarrow \uparrow$.

This result has been stated and proved in [10, Theorem 1] for separated bi-processes only. We are confident that it remains true for the more general bi-processes that we obtain after the transformations presented in this paper. To justify this claim, it is not feasible to adapt the proofs of the technical report associated to [10], which rely on a complex technical condition (called well-origination) that cannot be adapted without changing large proofs involving the inner workings of ProVerif's saturation procedure. However, the authors of [10] have indicated in personal communication that they hope to lift these conditions completely in order to obtain results that would be easier to reuse in future developments, but also to justify some recent features of ProVerif that would, in turn, make it easy to justify the soundness of the verification procedure on our extended bi-processes. More details are available in appendix A.

Without restrictions and when B is separated, the conclusion of Theorem 1 could be $\text{fst}(B) \approx_t \text{snd}(B)$. However, as already noted in [10], with restrictions, it does not directly show trace equivalence (or observational equivalence) for traces satisfying their respective restrictions (obtained through projections). We illustrate this observation below through an example.

Example 10: Going back to our running example, a natural way to encode unlinkability would be to consider the (bi)process P_{BH} as described in Example 3, and use a restriction to discard some traces on the right-hand side, enforcing that each tag executes only once. Applying this idea on the toy example developed in Example 8, we will consider the restriction ρ_2 which says that if two sessions of the same tag occur in a trace (on the right side of the bi-process), then this should be the same session. This is ensured by requiring that the session ids are the same, i.e. $x_2 = x'_2$.

When analysing an equivalence property, a restriction in ProVerif has to be seen as a bi-restriction, and will discard

bi-traces. Therefore, such a restriction will also discard some behaviours on the left-hand side. Thus, our example restriction ρ_2 actually does not encode the unlinkability property we want: Considering the process B given in Example 8, it is relatively easy to see that $\text{traces}(B)\downarrow\uparrow$ (both sides of the bi-process are equal up to a renaming). Thus, we also have convergence for the restricted set of bi-traces, i.e.

$$\text{traces}^{\rho_2}(B)\downarrow\uparrow$$

while our process clearly does not satisfy unlinkability.

Regarding Tamarin, restrictions and lemmas are something that have been available for quite a long time. However, the soundness of the tool in presence of restrictions for equivalence properties has only been established recently [26], and only for a specific type of restrictions which are actually not sufficient to write the restriction ρ_2 mentioned above.

C. Trace equivalence in presence of restrictions

In presence of restrictions, and when B is separated, we would like to be able to establish that $\text{traces}^{\mathcal{R}}(B)\downarrow\uparrow$ implies trace equivalence between $\text{fst}(B)$ and $\text{snd}(B)$ for traces satisfying their respective restrictions, i.e.

$$\text{traces}^{\text{fst}(\mathcal{R})}(\text{fst}(B)) \approx_t \text{traces}^{\text{snd}(\mathcal{R})}(\text{snd}(B)).$$

We first illustrate through examples why such a result does not hold in general.

Example 11: Consider the following bi-process

$$B = \text{event}(e(\text{diff}[\text{ok}^L, \text{ok}^R])); \text{out}(c, \text{diff}[\text{bad}^L, \text{bad}^R])$$

and the restriction

$$\rho = \text{event}(e(\text{diff}[x^L, x^R])) \Rightarrow x^L = x^R.$$

We have that $\text{traces}^{\rho}(B)$ contains only the empty bi-trace. However, $\text{fst}(\rho) \stackrel{\text{def}}{=} (\text{event}(e(x^L)) \Rightarrow x^L = x^R)$ will not discard any traces on $\text{fst}(B)$: recall that x^R is quantified existentially because it only occurs in the conclusion of the clause. The same goes for $\text{snd}(\rho)$ and $\text{snd}(B)$. Therefore, we have that:

$$\text{traces}^{\text{fst}(\rho)}(\text{fst}(B)) \not\approx_t \text{traces}^{\text{snd}(\rho)}(\text{snd}(B))$$

Note that the restriction ρ applies in a symmetric way on $\text{fst}(B)$ and $\text{snd}(B)$ (a variable occurring free on the right, as x^R , being quantified existentially). The issue here comes from the fact that this restriction discards more traces when considering the bi-process than its two projections individually. Concretely, a non-empty trace tr_1 of $\text{fst}(B)$ trivially satisfies the restriction $\text{fst}(\rho)$ but cannot be lifted to a bi-trace of B that satisfies ρ , i.e. there is no bi-trace $\text{tr} \in \text{traces}^{\rho}(B)$ such that $\text{fst}(\text{tr}) = \text{tr}_1$. The same issue holds for traces of $\text{snd}(B)$.

Example 12: Consider the bi-process

$$B = \text{event}(e(\text{diff}[\text{ok}^L, \text{ok}^R], \text{ok})); \text{out}(c, \text{ok})$$

and the following restriction:

$$\text{event}(e(\text{diff}[\text{ok}^R, \text{ok}^R], \text{diff}[x, x])) \Rightarrow x \neq \text{ok}$$

This restriction ρ will not discard any trace from $\text{traces}(B)$ as any event e occurring in $\text{traces}(B)$ will not have $\text{diff}[\text{ok}^R, \text{ok}^R]$ as first argument. Thus, we have that $\text{traces}^{\rho}(B)\downarrow\uparrow$, and it is possible to lift any trace possible on the left to a bi-trace T satisfying the bi-restrictions.

However, it is easy to see that the bi-trace

$$T = \text{event}(e(\text{diff}[\text{ok}^L, \text{ok}^R], \text{ok})); \text{out}(c, \text{ok})$$

is such that $\text{snd}(T) \not\vdash \text{snd}(\rho)$, whereas $\text{fst}(T) \vdash \text{fst}(\rho)$, thus we have that:

$$\text{traces}^{\text{fst}(\rho)}(\text{fst}(B)) \not\approx_t \text{traces}^{\text{snd}(\rho)}(\text{snd}(B))$$

In the remaining of this section, we establish the expected result (see Proposition 1) for the class of restrictions that are *diff-safe*. First, as shown in Example 11, it is important to be able to lift a trace from the left-hand side and satisfying $\text{fst}(\mathcal{R})$ to a bi-trace in $\text{traces}^{\mathcal{R}}(B)$. Second, as shown in Example 12, we have also to ensure that this bi-trace will satisfy the restrictions on the right. This leads us to consider the following definition.

Definition 6: Let $\mathcal{R}$ be a set of bi-restrictions, and B a separated closed bi-process. We say that $\mathcal{R}$ is *diff-safe* w.r.t. B if for all $T_{\text{fst}} \in \text{traces}^{\text{fst}(\mathcal{R})}(\text{fst}(B))$, there exists $T \in \text{traces}^{\mathcal{R}}(B)$ such that $\text{fst}(T) \cong T_{\text{fst}}$, and $\text{snd}(T) \vdash \text{snd}(\rho)$ for all $\rho \in \mathcal{R}$, and conversely for traces in $\text{traces}^{\text{snd}(\mathcal{R})}(\text{snd}(B))$.

We establish the following result.

Proposition 1: Let B be a separated closed bi-process, and $\mathcal{R}$ be a set of restrictions that is *diff-safe*. If $\text{traces}^{\mathcal{R}}(B)\downarrow\uparrow$ then

$$\text{traces}^{\text{fst}(\mathcal{R})}(\text{fst}(B)) \approx_t \text{traces}^{\text{snd}(\mathcal{R})}(\text{snd}(B)).$$

Proof. We establish the result for one direction. The other one can be obtained in a similar way. Let $T_{\text{fst}} \in \text{traces}^{\text{fst}(\mathcal{R})}(\text{fst}(B))$. We have to establish the existence of $T_{\text{snd}} \in \text{traces}^{\text{snd}(\mathcal{R})}(\text{snd}(B))$ such that $T_{\text{fst}} \cong T_{\text{snd}}$.

By definition of *diff-safe*, we know that there exists $T \in \text{traces}^{\mathcal{R}}(B)$ such that $\text{fst}(T) = T_{\text{fst}}$ and $\text{snd}(T) \vdash \text{snd}(\rho)$ for all $\rho \in \mathcal{R}$. Relying on the fact that $\text{traces}^{\mathcal{R}}(B)\downarrow\uparrow$, we deduce that $\text{fst}(T) \cong \text{snd}(T)$. As $T \in \text{traces}^{\mathcal{R}}(B)$, we know that $\text{snd}(T) \in \text{traces}(\text{snd}(B))$, and since $\text{snd}(T) \vdash \text{snd}(\rho)$ for all $\rho \in \mathcal{R}$, we deduce that $\text{snd}(T) \in \text{traces}^{\text{snd}(\mathcal{R})}(\text{snd}(B))$. $\square$

Deciding whether a set of restrictions is *diff-safe* is not an easy task. As a stepping stone towards a syntactic sufficient condition, we propose in the next lemma a useful characterisation that is used later on to identify some classes of bi-restrictions which are *diff-safe*.

Lemma 1: Let $\mathcal{R}$ be a set of bi-restrictions, and B be a separated closed bi-process such that $\text{traces}^{\mathcal{R}}(B)\downarrow\uparrow$, and such that for any $\rho \in \mathcal{R}$ and $T \in \text{traces}(B)$, we have:

$$T \vdash \rho \Leftrightarrow \text{fst}(T) \vdash \text{fst}(\rho) \Leftrightarrow \text{snd}(T) \vdash \text{snd}(\rho)$$

We have that $\mathcal{R}$ is *diff-safe* w.r.t. B .

Proof. Let $T_{\text{fst}} \in \text{traces}^{\text{fst}(\mathcal{R})}(\text{fst}(B))$. We have to establish that there exists $T \in \text{traces}^{\mathcal{R}}(B)$ such that $\text{fst}(T) \cong$

T_{fst} , and $\text{snd}(T) \vdash \text{snd}(\rho)$ for all $\rho \in \mathcal{R}$. We show this result by induction on the length of T_{fst} . First, the result trivially holds for a trace of length 0. We now assume that $T_{\text{fst}} = T'_{\text{fst}} \cdot \alpha$. Because the restriction satisfaction relation is prefix closed, i.e. $T_{\text{fst}} \vdash \rho$ implies $T'_{\text{fst}} \vdash \rho$, we have that $T'_{\text{fst}} \in \text{traces}^{\text{fst}(\mathcal{R})}(\text{fst}(B))$. Applying the induction hypothesis, we deduce that there exists $T' \in \text{traces}^{\mathcal{R}}(B)$ such that $\text{fst}(T') \cong T'_{\text{fst}}$, and $\text{snd}(T') \vdash \text{snd}(\rho)$ for all $\rho \in \mathcal{R}$. We know that $T' \in \text{traces}(B)$, and we distinguish two cases:

- We have that there exists β such that $T = T' \cdot \beta \in \text{traces}(B)$ with $\text{fst}(\beta) = \alpha$. In such a case, relying on our hypothesis, and since $\text{fst}(T) \vdash \text{fst}(\rho)$ for all $\rho \in \mathcal{R}$, we deduce that $T \vdash \rho$ for all $\rho \in \mathcal{R}$, and therefore $T \in \text{traces}^{\mathcal{R}}(B)$. We have also that $\text{fst}(T) \cong T_{\text{fst}}$, and $\text{snd}(T) \vdash \text{snd}(\rho)$ relying on our hypothesis. This allows us to conclude.
- We have that there does not exist β such that $T = T' \cdot \beta \in \text{traces}(B)$ and $\text{fst}(\beta) = \alpha$ meaning that $T' \in \text{traces}^{\mathcal{R}}(B)$ diverges which is not possible by hypothesis.

This concludes the proof of the first direction, and the other direction can be done in a similar way. $\square$

Example 13: Going back to Example 8, the restriction ρ_2 does not satisfy the criterion given in Lemma 1. It is indeed easy to see that for some $T \in \text{traces}(B)$, $\text{fst}(T)$ will satisfy $\text{fst}(\rho_2)$ whereas $\text{snd}(T)$ will not satisfy $\text{snd}(\rho_2)$. Actually, ρ_2 is not diff-safe. Consider the following bi-trace:

$$T' = \begin{cases} \text{event}(e(\text{sid}_1, \text{diff}[k^L, k^R])); \text{out}(c, \text{diff}[k^L, k^R]); \\ \text{event}(e(\text{sid}_2, \text{diff}[k^L, k^R])); \text{out}(c, \text{diff}[k^L, k^R]) \end{cases}$$

We have that $T_{\text{fst}} \in \text{traces}^{\text{fst}(\rho_2)}(B)$. The only bi-trace T_0 such that $T_0 \in \text{traces}^{\rho_2}(B)$ with $\text{fst}(T_0) \cong T_{\text{fst}}$ is actually T' . However, it happens that $\text{snd}(T') \not\vdash \rho_2$.

In ProVerif, restrictions are actually bi-restrictions, and as we have seen in Proposition 1, a link between bi-restrictions and their projections has to be established to get the expected equivalence, i.e.

$$\text{traces}^{\text{fst}(\mathcal{R})}(\text{fst}(B)) \approx_t \text{traces}^{\text{snd}(\mathcal{R})}(\text{snd}(B)).$$

This way of considering restrictions (i.e. referring to their projections) is more in line with what is done in Tamarin. Even if restrictions have been available for quite a long time in Tamarin, the soundness of the procedure in presence of restrictions has only been established recently, and for Type-0 restrictions only [26]. Intuitively, a Type-0 restriction is a Tamarin restriction that can be written relying solely on facts of arity 0. The idea behind this condition is to ensure that a restriction will apply in a symmetric way on both sides (on the dependency graph and also on its mirror). It may seem quite restrictive. However, it is important to note that, in Tamarin, any fact has a timestamp embedded in it. Therefore, Type-0 restrictions are expressive enough to enforce protocol phases or to model the fact that a certain action occurs only once:

$$\text{OnlyOnce}()@ \#i \ \& \ \text{OnlyOnce}()@ \#j \Rightarrow \#i = \#j$$

Restrictions having a similar effect can also be written in ProVerif but restrictions with events of arity 0 are not sufficient for that.

Example 14: Let B be a bi-process (with no replication) containing several actions, e.g. $\text{out}(c_1, u_1), \dots, \text{out}(c_n, u_n)$, and assume that we want to consider traces where at most one of these outputs is executed. The idea is to add `new sid; event(OnlyOnce(sid))` before each output, and to consider the following restriction:

$$\begin{aligned} &\text{event}(\text{OnlyOnce}(\text{diff}[x, x])) \wedge \\ &\text{event}(\text{OnlyOnce}(\text{diff}[y, y])) \Rightarrow x = y \end{aligned}$$

It is easy to see that such a restriction ρ satisfies the hypothesis given in Lemma 1, and is thus diff-safe, allowing one to derive the expected result, i.e. $\text{traces}^\rho(B) \downarrow \uparrow$ implies that $\text{traces}^{\text{fst}(\rho)}(\text{fst}(B)) \approx_t \text{traces}^{\text{snd}(\rho)}(\text{snd}(B))$. In other words, we get the soundness of ProVerif considering a more natural semantics for restrictions (we consider their respective projections).

Relying on diff-safe restrictions (or even simply the characterisation given in Lemma 1), we can also encode protocol phases (which is already a built-in mechanism in ProVerif). We believe that the criterion we provide with Lemma 1 to characterise diff-safe restrictions can be seen as the ProVerif counterpart of Type-0 restriction in the Tamarin setting. Thus, as done in [26] for Tamarin, we have also established the soundness of ProVerif w.r.t. diff-safe restrictions considering a more natural semantics for restrictions. Indeed, for this class of bi-restrictions, we can reason on their projections avoiding the issues mentioned in Examples 11 and 12.

However, the class of diff-safe restrictions is too restrictive to encode unlinkability as done e.g. in Example 10 (see Example 13). We thus develop another approach in the following section using the `allowDiffPatterns` option and axioms of ProVerif.

IV. OUR APPROACH

As we have seen in Example 7, ProVerif is not able to conclude on the bi-process modelling our unlinkability property on the Basic Hash protocol, and diff-safe restrictions are too restrictive to encode this property. We thus propose another approach. We propose to transform such bi-processes so that diff-equivalence may be achieved (and verified with ProVerif) on the transformed bi-processes, and that this diff-equivalence implies trace equivalence for the original process. We proceed in two steps:

- 1) We duplicate the get instructions to dissociate the process of the left and the one on the right.
- 2) We add some axioms to help ProVerif to reason on our bi-process. These axioms are used to improve the accuracy of the verification procedure by avoiding over-approximations.

Before formally defining these two steps in Sections IV-B and IV-C, we illustrate the approach on our running example.

A. Going back to our running example

On our running example, our transformations will only change the bi-process modelling the reader since it is the only one that features a get instruction. Instead of performing a get instruction to access one (bi)record in the keys table, it will perform two get instructions in a row to access two records in the *keys* table. This will allow one to choose two different records: one for the left and one for the right. Of course, the test performed remains the same and has to be satisfied by the element chosen on the left, and the one chosen on the right. We show the transformed bi-process in Figure 4, where the symbol $_$ represents irrelevant names for variables or public channels, and bad_L and bad_R are distinct public constants. The reader can ignore the added events at this stage.

This transformation alone is not enough to allow ProVerif to conclude. We thus add some axioms that emphasise some properties that are necessarily satisfied, to make sure that these properties are not lost in the over-approximations performed in ProVerif's verification procedure. To this end, we make use of the three events introduced in the transformed bi-process: Inserted, FailL, and FailR. The first one indicates when an element is inserted in the table *keys*, whereas the two others will be used to track the outcome of lookups in that table: FailL indicates a failure of the get instruction regarding the left hand side of the process, whereas FailR indicates a failure of the get instruction w.r.t. the right process. Relying on these events, we state two axioms given in Figure 4.

Let us focus on the first one. In case the FailL event is executed with a value M , it means that there is no entry $\text{diff}[N, N']$ in the table allowing one to evaluate the expression $\text{eq}(\text{proj}_2(M), \text{h}(\text{proj}_1(M), N))$ to true. Otherwise, we should have chosen this entry to execute the then branch. Thus, it is safe to say that for any entry $\text{diff}[N, N']$ in the table *keys*, we have $\text{proj}_2(M) \neq \text{h}(\text{proj}_1(M), N)$. The same applies on the right.

We may note that such an axiom is indeed satisfied since we know that insertions in the table can be done at the very beginning of the execution traces. This means that the element allowing one to perform the reasoning above can be assumed to be present before the execution of the FailL or FailR event. The validity of the axioms added by our transformation will be a consequence of Proposition 2.

B. Desynchronising the two parts of the bi-process

We now define the first step of the transformation informally described above.

Definition 7: Let B be a separated bi-process. We define $\mathcal{T}_1(B)$ as follows:

- $\mathcal{T}_1(\text{get } \text{tbl}(\text{diff}[x^L, x^R]) \text{ st. } D \text{ in } B_{\text{then}} \text{ else } B_{\text{else}})$ is the following bi-process, where $_$ represents irrelevant names for variables or public channels, and $\text{bad} =$

$\text{diff}[\text{bad}_L, \text{bad}_R]$ for two arbitrary distinct public constants bad_L and bad_R :

```

get tbl(diff[xL, _]) st. fst(D) in
  get tbl(diff[_ , xR]) st. snd(D) in  $\mathcal{T}_1(B_{\text{then}})$ 
  else out(_, bad)
else
  get tbl(diff[_ , xR]) st. snd(D) in out(_, bad)
  else  $\mathcal{T}_1(B_{\text{else}})$ 

```

- The transformation is homomorphic in all other cases: for instance, we define:
 $-\mathcal{T}_1(\text{out}(c, M); B') = \text{out}(c, M); \mathcal{T}_1(B')$, and
 $-\mathcal{T}_1(\text{in}(c, \text{diff}[x^L, x^R]); B') = \text{in}(c, \text{diff}[x^L, x^R]); \mathcal{T}_1(B')$.

Note that, because B is separated, $\text{fst}(D)$ cannot have x^R as a free variable, which would render the above transformation semantically dubious. In the above definition, when $\text{fst}(D)$ and $\text{snd}(D)$ disagree, we use the bi-process $\text{out}(_, \text{diff}[\text{bad}_L, \text{bad}_R])$ to signal that we want diff-equivalence to fail. To conclude in more cases, a variant could slightly relax this, under the assumption that B_{then} and B_{else} have the same structure. We do not encounter this situation in our case studies, and we left this extension as future work.

We establish the following result so that the “equivalence is true” returned by ProVerif on $\mathcal{T}_1(B_0)$ will allow one to conclude that trace equivalence holds between the two projections of B_0 , and thus our unlinkability property holds.

Theorem 2: Let B_0 be a separated closed bi-process such that $\mathcal{T}_1(B_0) \downarrow \uparrow$. We have that $\text{fst}(B_0) \approx_t \text{snd}(B_0)$.

This theorem will be proved relying on two lemmas that we now introduce. In these lemmas, when T is a bi-execution trace of $\mathcal{T}_1(B_0)$, we allow ourselves to write $\text{fst}(T)$ meaning that the first projection has been applied on each configuration of T . Note that this does not imply that $\text{fst}(T)$ is an execution of $\text{fst}(\mathcal{T}_1(B_0))$ as $\mathcal{T}_1(B_0)$ is not separated. Actually, we only use $\text{fst}(T)$ to assess its indistinguishability w.r.t. another trace (e.g., in $T_{\text{fst}} \cong \text{fst}(T)$): we thus only care about the observables, and the first projection of the last frame of T . Since $\mathcal{T}_1(B_0)$ duplicates the get instruction, to establish the connection between $\text{fst}(B_0)$ and $\mathcal{T}_1(B_0)$, it is easier to consider a notion of execution trace where the two get instructions corresponding to the same get instruction in the original bi-process B_0 are triggered in a row. We denote $\overline{\text{traces}}$ this notion of bi-execution trace.

Lemma 2: Let B_0 be a separated closed bi-process such that $\mathcal{T}_1(B_0) \downarrow \uparrow$, and $T_{\text{fst}} \in \overline{\text{traces}}(\text{fst}(B_0))$. There exists $T \in \overline{\text{traces}}(\mathcal{T}_1(B_0))$ such that $T_{\text{fst}} \cong \text{fst}(T)$. A similar result holds for snd .

Proof. (Sketch) We will find T such that $\text{fst}(T)$ and T_{fst} differ only on the duplicated τ actions corresponding to transformed get operations. Although the lemma is stated in terms of static equivalence, frames are actually equal. The convergence hypothesis guarantees that we avoid $\text{diff}[\text{bad}_L, \text{bad}_R]$ which can artificially appear in the traces of the transformed bi-process. $\square$

```

in(c, diff[xL, xR]);
get keys(diff[yL, _]) st. eq(proj2(xL), h(proj1(xL, yL))) in
  get keys(diff[_ , yR]) st. eq(proj2(xR), h(proj1(xR, yR))) in
    out(c, ok)
  else event(FailR(xR)); out(_, diff[badL, badR])
else
  event(FailL(xL));
  get keys(diff[_ , yR]) st. eq(proj2(xR), h(proj1(xR, yR))) in
    out(_, diff[badL, badR])
  else out(c, error)

event(FailL(xL)) ∧ event(Inserted(diff[yL, yR])) ⇒ proj2(xL) ≠ h(proj1(xL, yL))
event(FailR(xR)) ∧ event(Inserted(diff[yL, yR])) ⇒ proj2(xR) ≠ h(proj1(xR, yR))

```

Fig. 4. Transformed reader bi-process for Basic Hash

Lemma 3: Let B_0 be a separated closed bi-process such that $\mathcal{T}_1(B_0) \downarrow$, and $T \in \overline{\text{traces}}(\mathcal{T}_1(B_0))$. There exists $T_{\text{fst}} \in \text{traces}(\text{fst}(B_0))$ such that $T_{\text{fst}} \cong \text{fst}(T)$. A similar result holds for snd .

Proof. (Sketch) The projected traces are simply obtained by de-duplicating the τ actions corresponding to `get` sub-processes. Again, the convergence hypothesis guarantees that we avoid `diff[badL, badR]`. $\square$

C. Refining the analysis in failure branches

We transform a *model*, i.e. a bi-process together with restrictions, axioms and lemmas, to a new model that is easier to verify by ProVerif. We first illustrate this transformation on a very simple toy example.

Example 15: Consider the following bi-process B :

```

insert tbl(ok);
get tbl(x) st. true in out(c, ok)
  else out(c, diff[okL, okR])

```

On this simple example, ProVerif returns “observational equivalence cannot be proved”, whereas it is relatively easy to see that the `else` branch cannot be taken (as there is an element in the table), and thus equivalence holds. Applying the transformation that we will formally present next will lead to the following bi-process:

```

event(Inserted(ok)); insert tbl(ok);
get tbl(x) st. true in out(c, ok)
  else event(Fail()); out(c, diff[okL, okR])

```

together with the following axiom:

$$\text{event(Fail())} \wedge \text{event(Inserted(diff[y^L, y^R]))} \Rightarrow \text{false}.$$

On this model, ProVerif is able to conclude that equivalence holds.

We now formally define this transformation.

Definition 8: We have that

$$(B, \mathcal{R}, \mathcal{A}x, \mathcal{L}) \rightsquigarrow_{\text{else}} (B', \mathcal{R}, \mathcal{A}x \cup \{ax\}, \mathcal{L})$$

when there exist a context C , two terms M_1 and M_2 without `diff` operators, and two processes P and Q such that:

- B is of the form:

$$C[\text{get tbl(diff[x^L, x^R]) st. eq(M₁, M₂) in } P \text{ else } Q]$$

- B' is of the form:

$$C[\text{get tbl(diff[x^L, x^R]) st. eq(M₁, M₂) in } P \text{ else event(Fail}(\vec{n}, \vec{y})) \cdot Q]$$

where $\vec{n} = \text{names}(\{M_1, M_2\})$, $\vec{y} = \text{vars}(\{M_1, M_2\}) \setminus \{x^L, x^R\}$, and all operations `insert tbl(M)` have been replaced by `event(Inserted(M)).insert tbl(M)`

- ax is of the form:

$$\text{event(Fail}(\vec{z}, \vec{y})) \wedge \text{event(Inserted(diff[x^L, x^R]))} \Rightarrow (M_1 \neq M_2) \{ \vec{n} \mapsto \vec{z} \}$$

where `Fail` and `Inserted` are distinct events that do not occur in $(B, \mathcal{R}, \mathcal{A}x, \mathcal{L})$, and $\vec{z}$ is a vector of fresh variables of the same length as $\vec{n}$.

Intuitively, the transformation instruments the process with events in order to express, in the added axiom, that we can only execute the `else` branch of the bi-process if no value inserted in `tbl` satisfies $M_1 = M_2$. This makes sense even if x^L and x^R do not occur in $\text{eq}(M_1, M_2)$: in that case we are simply saying that we can only visit the `else` branch if either $\text{eq}(M_1, M_2)$ is false, or no value has even been inserted in the table.

The above intuition is not entirely exact, though. Indeed, the semantics of the axiom is that if at some point in the trace the events `Fail` and `Inserted` have happened (for some values of the free variables $\vec{z}, \vec{y}, x^L, x^R$), then $\text{eq}(M_1, M_2)$ evaluates to true. The semantics of bi-processes enforces a weaker version of this statement, where the `Inserted` must happen before `Fail`. It is not possible to express such precise temporal properties in axioms with ProVerif but, as we shall see, we can live with this limitation.

The transformation can be iterated. Obviously, the `insert` events of several transformations relying on the same table can be merged. We justify next a single step of the transformation; the result extends naturally for iterations.

Proposition 2: Let $(B, \mathcal{R}, \mathcal{A}x, \mathcal{L})$, and $(B', \mathcal{R}, \mathcal{A}x', \mathcal{L})$ be two models such that $(B, \mathcal{R}, \mathcal{A}x, \mathcal{L}) \rightsquigarrow_{\text{else}} (B', \mathcal{R}, \mathcal{A}x', \mathcal{L})$. We have that $\mathcal{A}x' = \mathcal{A}x \cup \{ax\}$ for some ax . Assuming that all insert instructions are performed in a phase that precedes the phase of the modified get instruction, we have:

$$\text{traces}^{\mathcal{R}}(B') \vdash ax$$

Proof. Consider a bi-trace $T' \in \text{traces}^{\mathcal{R}}(B')$, and assume that $T' \not\vdash ax$. Then, for some substitution θ , we have $\text{Inserted}(\text{diff}[x^L, x^R])\theta$ in the bi-trace, as well as $\text{Fail}(\bar{z}, \bar{y})\theta$, and $(M_1 =_E M_2)\{\bar{n} \mapsto \bar{z}\}\theta$ is true. In other words, the execution of the get action, for the values given by θ , has stepped to the else branch, when there existed a bi-term $\text{diff}[x^L, x^R]\theta$ in the table that satisfies the condition. This contradicts the semantics of bi-processes. $\square$

This result, together with Theorem 2 and Theorem 1, shows that if ProVerif concludes on the transformed model, then trace equivalence holds for the original one. More precisely, we have the following result.

Theorem 3: Let $(B_0, \emptyset, \mathcal{A}x, \mathcal{L})$ be a separated ProVerif model, and $(B', \emptyset, \mathcal{A}x', \mathcal{L})$ be the model obtained after applying our two transformation, i.e. a model such that $(\mathcal{T}_1(B_0), \emptyset, \mathcal{A}x, \mathcal{L}) \rightsquigarrow_{\text{else}}^* (B', \emptyset, \mathcal{A}x', \mathcal{L})$. Moreover, we assume that:

- for all $\varrho \in \mathcal{A}x$, we have that $\text{traces}(B_0) \vdash \varrho$;
- for all $\varrho \in \mathcal{A}x$, we have that $\text{traces}(\mathcal{T}_1(B_0)) \vdash \varrho$;
- all insert instructions in B_0 are performed in a phase that precedes the phase of the get instructions transformed by $\rightsquigarrow_{\text{else}}$ above; and
- ProVerif establishes equivalence on the resulting model $(B', \emptyset, \mathcal{A}x', \mathcal{L})$ i.e. returns “diff-equivalence is true” on this model.

We conclude that $\text{fst}(B_0) \approx_t \text{snd}(B_0)$.

Note that the two first items are there to justify the axioms already present in the original model (not the ones added by our transformation).

Proof. By Proposition 2, we have $\text{traces}(B') \vdash \varrho'$ for any $\varrho' \in \mathcal{A}x' \setminus \mathcal{A}x$. By hypothesis, we have $\text{traces}(\mathcal{T}_1(B_0)) \vdash \varrho$ for any $\varrho \in \mathcal{A}x$, and this is also true for B' since the events inserted in the last transformation are new. Moreover, we have no restriction here. By Theorem 1, we deduce that $\text{traces}(B') \downarrow \uparrow$, and we also have $\text{traces}(\mathcal{T}_1(B_0)) \downarrow \uparrow$ since B' and $\mathcal{T}_1(B_0)$ are the same up to some events. We finally obtain $\text{fst}(B_0) \approx_t \text{snd}(B_0)$ by Theorem 2. $\square$

V. CASE STUDIES

In this section, we explain how our approach can be used to analyse unlinkability on several existing protocols.

A. Implementation

In order to analyse our protocols of interest, we have developed a tool that implements the two transformations described in the previous section. This tool makes it easy to test our approach on various protocols, and avoids mistakes during transformations. It takes as input a ProVerif model containing

a bi-process specifying the equivalence to be verified and returns another ProVerif file. If ProVerif concludes that diff-equivalence holds on the output file then our result allows one to conclude that the left and right processes written in the input file are in trace equivalence.

In practice, we implemented a ProVerif front-end tool to automatically apply our two transformations ($\approx 2k$ OCaml LoC). It takes as input a description of the protocol as presented in Example 7 and outputs a new ProVerif model in which the two parts of the bi-process have been desynchronised and the axioms corresponding to else branches have been added. ProVerif can then be ran on this transformed model to automatically prove the security of the protocol. The source code of our tool and material to reproduce results can be found at [19]. This includes ProVerif models for our case studies but also the generated files corresponding to their transformations: the reader may thus inspect the effect of our transformation without running our tool.

B. Results

We first study four examples: Basic Hash (used as a running example), Hash-Lock, Feldhofer, and a variant of the LAK protocol (LAK-v1). All these examples are described, e.g., in [2]. They are similar in the sense that they involve monotonic states, i.e. states that are never updated. On all these examples, it was already well-known that ProVerif is unable to conclude on the model encoding unlinkability relying on standard bi-process and diff-equivalence. Using our tool, we obtain a file on which ProVerif is able to conclude in a few seconds, thus we conclude that these protocols ensure unlinkability by Theorem 3. A detailed description of how our transformations work on the Basic Hash protocol is given in Section IV-A.

We then consider a fixed version of the OSK protocol, described as OSK-v2 in [2]. The protocol can be described informally as follows:

$$\begin{aligned} T \rightarrow R : & \quad g(h(k_T)) \\ T \text{ updates } & k_T \text{ with } h(k_T) \end{aligned}$$

Here, g and h are two hash functions. Each tag has a secret key k_T , whose initial value is stored in the reader’s database. The tag updates its key by applying the h function symbol at each session on its current key. The reader expects a message of the form $g(h^n(x))$ for some database entry x (h^n means that h is applied n times). This check is modelled relying on a private destructor $\text{retrievekey}(\cdot)$ allowing one to extract k from a term of the form $h(n, k)$ used to model $h^n(k)$. As usual in ProVerif, we encode the state of the tag using a private channel. To avoid false attacks, we rely on some axioms to improve precision in the manipulation of private channels. This is done following the approach developed in GSVerif [14] for trace-based properties. Extending the approach of [14] for equivalence properties is beyond the scope of this paper, but on this particular example, it is easy to be convinced that these axioms are valid, and thus our encoding correct. Despite this and the addition of a lemma to help regarding termination,

ProVerif is not able to establish that the equivalence holds. However, ProVerif concludes successfully on the transformed model (produced by our tool) hence unlinkability holds for OSK-v2 by Theorem 3.

C. Comparison with earlier work

In [2], Baelde *et al.* develop an approach to tackle the same class of protocols as the one we consider here, and using the same definition for unlinkability. It is thus natural to compare in more details this work with the present one.

First, we note that all the results presented above are consistent with the security analyses performed in [2] where unlinkability has been established relying on three sufficient conditions proved using the tool Tamarin.

Second, the verification technique proposed by [2], based on sufficient conditions, is at a high-level independent of the verification tool, whereas our work is specific to ProVerif. Moreover, the sufficient conditions themselves can provide some insight as to what makes unlinkability hold, or fail.

However, in practice, verifying the sufficient conditions of [2] has proved very difficult, and the authors of that paper could only obtain semi-automated Tamarin proofs. In general, we thus reckon that this verification technique will require more effort than our direct approach. In each of the case studies of [2], two different models had to be designed to verify the three sufficient conditions. Each time, several lemmas have to be written corresponding to the various conditions (between two and four depending on the protocol) and, in some cases, several intermediate lemmas are needed to help Tamarin prove the conditions.

In particular, regarding OSK-v2, the security analysis conducted in [2] has only been possible thanks to several simplifications and using some modelling tricks. For instance, one of the sufficient conditions has been established discarding the reader role, and the attacker was not given the power to apply a hash function on some terms.

Overall, our approach brings an interesting complement to the technique of [2], which can significantly ease the verification effort when ProVerif is applicable.

VI. CONCLUSION

We have shown that unlinkability can be automatically proved in ProVerif, by applying two generic transformations on the standard bi-process expressing unlinkability. Our approach significantly improves over existing automated proofs, avoiding several over-approximations as well as manual work. It builds on several specific features of ProVerif. In particular, we make crucial use of extended bi-processes with diff operators on bound variables, that have seen only limited use so far (and which would not make sense in e.g. Tamarin). More generally, we make use of the recent addition of axioms, lemmas and restrictions in ProVerif, and we provide a better understanding of how restrictions can be used to prove equivalences.

It would be natural to extend our theoretical development as well as our tool to generalise our second transformation

beyond equality tests between constructor terms. To deal with arbitrary terms D , a more involved transformation which essentially consists in computing the variants of an expression to get rid of destructor symbols will be necessary [16]. In light of the usefulness and simplicity of our approach, a promising perspective would be to integrate it by default in ProVerif, but this requires the addition of temporal constraints in correspondence queries. Meanwhile, this work introduces a new approach which is not specific to unlinkability. We develop in [19] an example (actually a toy protocol provided in [23]) which demonstrates that our approach can be used to prove *anonymity* properties. We think that a similar approach could have been used in [7] to overcome ProVerif limitations when manipulating lookup tables for anonymity properties. It appears as an alternative to the encoding used in [7] (and detailed in [6] - Section 5.4.3). Their encoding is based on the use of restrictions which have to be manipulated with a lot of care when analysing equivalence-based properties.

Our contribution in this paper, namely relaxing diff-equivalence to be able to prove unlinkability in ProVerif, does not address other difficult aspects that might arise in this task. Handling the tag's mutable state in the OSK-v2 case study has proved difficult, but there is hope that advances with GSVerif will soon solve this issue. It is also natural to consider protocols with state updates on the reader's side, i.e. updates of database entries. Although such operations are not directly supported in ProVerif, there is hope to model them naturally enough using restrictions – which calls for a (slight) generalisation of our results.

REFERENCES

- [1] M. Arapinis, T. Chothia, E. Ritter, and M. Ryan. Analysing unlinkability and anonymity using the applied pi calculus. In *Proc. 23rd IEEE Computer Security Foundations Symposium, CSF 2010, Edinburgh, United Kingdom, July 17-19, 2010*, pages 107–121. IEEE Computer Society, 2010.
- [2] D. Baelde, S. Delaune, and S. Moreau. A method for proving unlinkability of stateful protocols. In *Proc. of the 33rd IEEE Computer Security Foundations Symposium (CSF'20)*. IEEE Computer Society Press, July 2020.
- [3] D. Basin, R. Sasse, and J. Toro-Pozo. The EMV standard: Break, fix, verify. In *2021 IEEE Symposium on Security and Privacy (S&P'21)*, pages 1766–1781. IEEE, 2021.
- [4] D. A. Basin, J. Dreier, and R. Sasse. Automated symbolic proofs of observational equivalence. In I. Ray, N. Li, and C. Kruegel, editors, *Proc. 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS'15)*, pages 1144–1155. ACM, 2015.
- [5] K. Bhargavan, B. Blanchet, and N. Kobeissi. Verified models and reference implementations for the TLS 1.3 standard candidate. In *2017 IEEE Symposium on Security and Privacy, (S&P'17), San Jose, CA, USA, May 22-26, 2017*, pages 483–502. IEEE Computer Society, 2017.
- [6] K. Bhargavan, V. Cheval, and C. Wood. Handshake Privacy for TLS 1.3 - Technical report. Research report, Inria Paris ; Cloudflare, Mar. 2022.
- [7] K. Bhargavan, V. Cheval, and C. Wood. A symbolic analysis of privacy for tls 1.3 with encrypted client hello. In *Proceedings of the 29th ACM Conference on Computer and Communications Security (CCS'22)*, Los Angeles, USA, Nov. 2022. ACM Press.
- [8] B. Blanchet. An efficient cryptographic protocol verifier based on prolog rules. In *14th IEEE Computer Security Foundations Workshop (CSFW'01), 11-13 June 2001, Cape Breton, Nova Scotia, Canada*, pages 82–96, 2001.
- [9] B. Blanchet, M. Abadi, and C. Fournet. Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming*, 75(1):3–51, 2008.

- [10] B. Blanchet, V. Cheval, and V. Cortier. Proverif with lemmas, induction, fast subsumption, and much more. In *Proceedings of the 42nd IEEE Symposium on Security and Privacy (S&P'22)*. IEEE Computer Society Press, 2022.
- [11] M. Brusò, K. Chatzikokolakis, and J. den Hartog. Formal verification of privacy for RFID systems. In *Proc. 23rd IEEE Computer Security Foundations Symposium, (CSF'10)*, pages 75–88. IEEE Computer Society, 2010.
- [12] M. Brusò, K. Chatzikokolakis, S. Etalle, and J. den Hartog. Linking unlinkability. In C. Palamidessi and M. D. Ryan, editors, *Trustworthy Global Computing - 7th International Symposium, TGC 2012, Newcastle upon Tyne, UK, September 7-8, 2012, Revised Selected Papers*, volume 8191 of *LNCS*, pages 129–144. Springer, 2012.
- [13] V. Cheval, V. Cortier, and S. Delaune. Deciding equivalence-based properties using constraint solving. *Theor. Comput. Sci.*, 492:1–39, 2013.
- [14] V. Cheval, V. Cortier, and M. Turuani. A little more conversation, a little less action, a lot more satisfaction: Global states in proverif. In *Proceedings of the 31st IEEE Computer Security Foundations Symposium (CSF'18)*, pages 344–358, 2018.
- [15] T. Chothia and V. Smirnov. A traceability attack against e-passports. In R. Sion, editor, *Financial Cryptography and Data Security, 14th International Conference, FC 2010, Tenerife, Canary Islands, Spain, January 25-28, 2010, Revised Selected Papers*, volume 6052 of *LNCS*, pages 20–34. Springer, 2010.
- [16] H. Comon-Lundh and S. Delaune. The finite variant property: How to get rid of some algebraic properties. In *Proceedings of the 16th International Conference on Rewriting Techniques and Applications (RTA'05)*, volume 3467 of *LNCS*, pages 294–307, Nara, Japan, Apr. 2005. Springer.
- [17] V. Cortier, A. Dallon, and S. Delaune. Efficiently deciding equivalence for standard primitives and phases. In J. López, J. Zhou, and M. Soriano, editors, *Proc. 23rd European Symposium on Research in Computer Security, (ESORICS'18)*, volume 11098 of *LNCS*, pages 491–511. Springer, 2018.
- [18] C. Cremers, M. Horvat, J. Hoyland, S. Scott, and T. van der Merwe. A comprehensive symbolic analysis of TLS 1.3. In B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, editors, *Proc. ACM SIGSAC Conference on Computer and Communications Security, (CCS'17)*, pages 1773–1788. ACM, 2017.
- [19] A. Debant, S. Delaune, and D. Baelde. Proving Unlinkability using ProVerif through Desynchronized Bi-Processes. working paper or preprint available at <https://hal.inria.fr/hal-03674979>, May 2022.
- [20] J. Dreier, L. Hirschi, S. Radomirovic, and R. Sasse. Automated unbounded verification of stateful cryptographic protocols with exclusive OR. In *Proc. 31st IEEE Computer Security Foundations Symposium (CSF'18)*, pages 359–373. IEEE Computer Society, 2018.
- [21] I. Filimonov, R. Horne, S. Mauw, and Z. Smith. Breaking unlinkability of the ICAO 9303 standard for e-passports using bisimilarity. In K. Sako, S. A. Schneider, and P. Y. A. Ryan, editors, *Computer Security - ESORICS 2019 - 24th European Symposium on Research in Computer Security, Luxembourg, September 23-27, 2019, Proceedings, Part I*, volume 11735 of *Lecture Notes in Computer Science*, pages 577–594. Springer, 2019.
- [22] L. Hirschi, D. Baelde, and S. Delaune. A method for verifying privacy-type properties: The unbounded case. In *IEEE Symposium on Security and Privacy, (S&P'16), San Jose, CA, USA, May 22-26, 2016*, pages 564–581. IEEE Computer Society, 2016.
- [23] L. Hirschi, D. Baelde, and S. Delaune. A method for unbounded verification of privacy-type properties. *Journal of Computer Security*, 27(3):277–342, 2019.
- [24] R. J. Horne, S. Mauw, and S. Yurkov. Unlinkability of an improved key agreement protocol for EMV 2nd gen payments. In *Proceedings of the 35th IEEE Computer Security Foundations Symposium (CSF'22)*, Haifa, Israel, Aug. 2022. IEEE Computer Society Press.
- [25] S. Meier, B. Schmidt, C. Cremers, and D. A. Basin. The TAMARIN prover for the symbolic analysis of security protocols. In *Proc. 25th International Conference on Computer Aided Verification (CAV'13)*, pages 696–701, 2013.
- [26] P. Paradzik and A. Derek. Conditional observational equivalence and off-line guessing attacks in multiset rewriting. In *Proceedings of the 35th IEEE Computer Security Foundations Symposium (CSF'22)*. IEEE Computer Society, 2022.
- [27] T. van Deursen, S. Mauw, and S. Radomirovic. Untraceability of RFID protocols. In *Information Security Theory and Practices. Smart Devices,*

Convergence and Next Generation Networks, Second IFIP WG 11.2 International Workshop, WISTP 2008, Seville, Spain, May 13-16, 2008. Proceedings, volume 5019 of *LNCS*, pages 1–15. Springer, 2008.

APPENDIX

A. On the soundness of ProVerif for non-separated bi-processes

Theorem 1 is established in [10] for separated bi-processes. We outline the argument next, referring to results in the technical report associated to that paper. In order to verify diff-equivalence, ProVerif first translates the definition of cryptographic primitives and the initial bi-process into a set of Horn clauses. It is shown in [10, Lemma 11] that if the bi-process admits non-convergent traces, then there exists a logical derivation of some bad atom using the Horn clauses. Then, [10, Theorem 3] justifies that the derivability of bad would be detected by the saturation step of ProVerif's algorithm.

The first step in this reasoning easily carries over to non-separated bi-processes. In fact, the generalised bi-processes made available by the `allowDiffPatterns` option are nothing more than a back-translation: mechanisms already available in ProVerif clauses are added to bi-processes, with a semantics that naturally follows the Horn clause semantics. We briefly explain this point below, referring the reader to [23, §5.1] for a more detailed discussion, and to [10, §4.3.2] for a full definition of the translation.

a) *Translating bi-processes into Horn clauses:* Clauses in ProVerif use an atomic proposition $\text{att}'(m^L, m^R)$ to express the fact that the attacker can compute m^L using some recipe after executing some trace with the left projection of the bi-process, and would compute m^R using the same recipe after executing the same trace against the right projection. In terms of bi-executions, this simply means that some recipe yields $\text{diff}[m^L, m^R]$ when applied to the bi-frame obtained after executing the trace on the bi-process.

Consider for example the following separated bi-process, where `cst` is an arbitrary constant:

$$\text{in}(c, x); \text{if } x = \text{cst} \text{ then } \text{out}(c, x) \text{ else } \text{out}(c, m)$$

The translation would notably produce the following clauses for that bi-process (we assume that m does not contain occurrences of x , but may contain diff operators):

$$\begin{aligned} \text{att}'(\text{cst}, \text{cst}) &\Rightarrow \text{att}'(\text{cst}, \text{cst}) \\ \text{att}'(x^L, x^R) \wedge x^L \neq \text{cst} \wedge x^R \neq \text{cst} &\Rightarrow \text{att}'(\text{fst}(m), \text{snd}(m)) \\ \text{att}'(\text{cst}, x^R) \wedge x^R \neq \text{cst} &\Rightarrow \text{bad} \end{aligned}$$

The first clause encodes the execution of our bi-process when the conditional succeeds on both projections; it brings no new knowledge to the attacker. The second one corresponds to the case where the conditional fails in both projections. The last one corresponds to a non-convergence case, where the conditional passes on the left but not on the right.

As is visible in this example, input variables are duplicated during the translation into their left and right counterparts. We

can back-translate this to bi-processes, obtaining for instance a bi-process that is rigorously equivalent to the previous one for ProVerif's translation (and which we make equivalent as a bi-process by providing it with the suitable semantics):

```
in(c, diff[xL, xR]);
if diff[xL, xR] = cst then out(c, diff[xL, xR])
else out(c, m)
```

More interestingly, we can modify our process to only perform the test on the left projection, as we do in our first transformation:

```
in(c, diff[xL, xR]);
if xL = cst then out(c, diff[xL, xR])
else out(c, m)
```

The generated clauses would now become:

$$\begin{aligned} \text{att}'(\text{cst}, x^R) &\Rightarrow \text{att}'(\text{cst}, x^R) \\ \text{att}'(x^L, x^R) \wedge x^L \neq \text{cst} &\Rightarrow \text{att}'(\text{fst}(m), \text{snd}(m)) \end{aligned}$$

Note that the clause deriving bad has disappeared since the test $x^L = \text{cst}$ cannot succeed on one projection only, since it does not contain diff operators anymore. Also note that the second clause above is simply obtained from the second clause of the first bi-process by dropping the atom corresponding to the right projection of the original conditional.

The fact that [10, Lemma 1] remains true for extended bi-processes simply stems from the fact that the semantics of extended bi-processes (as presented in this paper) are adequately captured by this simple extension of the translation process, which is the one currently implemented in ProVerif.

b) Soundness of the saturation procedure: As it is, [10, Theorem 3] relies on the assumption that clauses satisfy a so-called well-origination condition. The condition notably imposes that, if a clause concludes with $\text{att}'(m, n)$, then the variables in m (resp. n) are introduced in some hypothesis of the clause. The definition is actually unclear in the case of clauses corresponding to biprocesses and, due to the complexity of the proofs, we could not determine how to clarify it.

The authors of [10] have indicated in personal communication that they expect to lift this technical condition completely in the future, which would almost immediately justify our use of extended biprocesses. Their goal is obviously to clarify and generalise the result, but also to justify an option that is currently available in ProVerif (under the name `maxHyps`) which allows the saturation procedure to drop hypotheses in large clauses in order to achieve termination more easily. This perspective would provide much more reasonable ways to justify ProVerif's soundness when used on (a class of) non-separated bi-processes. In fact, the justification of `maxHyps` would automatically allow the clauses resulting from the bi-processes used in our transformations: as in the example above, our bi-processes only exploit "unilateral conditionals", and their clauses are obtained from the clauses of standard bi-processes by dropping the unwanted projection of some conditionals.

B. Proofs of Section IV

Given a multiset $\mathcal{P}$ of bi-processes, we define $\mathcal{T}_1(\mathcal{P})$ the multiset of all transformed elements of $\mathcal{P}$. We extend this notion to bi-configurations as follows: $\mathcal{T}_1(K) = (\mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i)$ when $K = (\mathcal{P}; \Phi; \mathcal{S}; i)$. Similarly, we extend the notion `fst` and `snd` to multiset of processes and to configurations as follows: $\text{fst}(K) = (\text{fst}(\mathcal{P}); \Phi_{\text{fst}}; \mathcal{S}_{\text{fst}}; i)$ when $K = (\mathcal{P}; \Phi; \mathcal{S}; i)$, and similarly for `snd`.

Lemma 2: Let B_0 be a separated closed bi-process such that $\mathcal{T}_1(B_0) \downarrow \uparrow$, and $T_{\text{fst}} \in \text{traces}(\text{fst}(B_0))$. There exists $T \in \text{traces}(\mathcal{T}_1(B_0))$ such that $T_{\text{fst}} \cong \text{fst}(T)$. A similar result holds for `snd`.

Proof. We prove the result for `fst`, the case of `snd` being symmetric. We start by establishing the following claim.

Claim. Let K_0 be a separated bi-configuration, and $K = \mathcal{T}_1(K_0)$ such that $K \downarrow \uparrow$. If $\text{fst}(K_0) \xrightarrow{\alpha} K'_L$ for some α and some K'_L , then there exist a separated bi-configuration K'_0 and K' such that $K' \downarrow \uparrow$, $K' = \mathcal{T}_1(K'_0)$, and $\text{fst}(K'_0) = K'_L$. Moreover, either this step corresponds to the execution of a `get` instruction, and we have that $\alpha = \tau$, and $K \xrightarrow{\tau} K'$ (with the execution of two `get` instructions in a row); otherwise we have that $K \xrightarrow{\alpha} K'$.

Note that the result stated above allows us to prove the statement in the lemma by induction on the length of the execution trace T_{fst} . Now, it remains to establish this claim. We do a case analysis depending on the rule used to perform $\text{fst}(K_0) \xrightarrow{\alpha} K'_L$.

Case IN. In such a case, we have that:

- $K_0 = (\{i : \text{in}(c, \text{diff}[x^L, x^R]); P\} \cup \mathcal{P}; \Phi; \mathcal{S}; i);$
- $\alpha = \text{in}(c, R)$ for some recipe R such that $R\Phi_{\text{fst}} \downarrow =_E M$, and $R\Phi_{\text{snd}} \downarrow =_E M'$;
- $K'_L = (\{i : \text{fst}(P)\{x^L \mapsto M\}\} \cup \text{fst}(\mathcal{P}); \Phi_{\text{fst}}; \mathcal{S}_{\text{fst}}; i);$ and
- $K = \mathcal{T}_1(K_0) = (\{i : \text{in}(c, \text{diff}[x^L, x^R]); \mathcal{T}_1(P)\} \cup \mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i).$

We consider the two following bi-configurations:

- $K'_0 = (\{i : P\{x^L \mapsto M, x^R \mapsto M'\}\} \cup \mathcal{P}; \Phi; \mathcal{S}; i);$ and
- $K' = (\{i : \mathcal{T}_1(P)\{x^L \mapsto M, x^R \mapsto M'\}\} \cup \mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i).$

Since K_0 is separated, we have that K'_0 is separated too, and we have that $\text{fst}(K'_0) = K'_L$. Relying on the fact that our transformation is compatible with the substitution of diff-free terms, i.e., $\mathcal{T}(B\sigma) = \mathcal{T}(B)\sigma$ for all bi-process B and substitution σ which do not introduce terms with diff operators, we deduce that $\mathcal{T}_1(K'_0) = K'$. Then, it is easy to see that $K \xrightarrow{\text{in}(c, R)} K'$, and thus $K' \downarrow \uparrow$.

Most of the cases can be done in a similar way. We focus now on the more interesting cases, i.e. the ones corresponding to `get` instructions.

Case GET-T. In such a case, we have that $\alpha = \tau$, and:

- $K_0 = (\{i : \text{get } \text{tbl}(\text{diff}[x^L, x^R]) \text{ st. } D \text{ in } B_{\text{then}} \text{ else } B_{\text{else}}\} \cup \mathcal{P}; \Phi; \mathcal{S}; i);$

- $K'_L = (\{i : \text{fst}(B_{\text{then}})\{x^L \mapsto M_L\}\} \cup \text{fst}(\mathcal{P}); \Phi_{\text{fst}}; \mathcal{S}_{\text{fst}}; i)$ where M_L is such that $\text{tbl}(M_L) \in \mathcal{S}_{\text{fst}}$, and $\text{fst}(D)\{x^L \mapsto M_L\} \Downarrow =_{\text{E}} \text{true}$;
- $K = \mathcal{T}_1(K_0) = (\{i : B'\} \cup \mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i)$ where:

$$B' = \begin{cases} \text{get } \text{tbl}(\text{diff}[x^L, _]) \text{ st. } \text{fst}(D) \text{ in} \\ \quad \text{get } \text{tbl}(\text{diff}[_, x^R]) \text{ st. } \text{snd}(D) \text{ in} \\ \quad \mathcal{T}_1(B_{\text{then}}) \\ \text{else out}(_, \text{diff}[\text{bad}_L, \text{bad}_R]) \\ \text{else} \\ \quad \text{get } \text{tbl}(\text{diff}[_, x^R]) \text{ st. } \text{snd}(D) \text{ in} \\ \quad \text{out}(_, \text{diff}[\text{bad}_L, \text{bad}_R]) \\ \text{else } \mathcal{T}_1(B_{\text{else}}) \end{cases}$$

Let M_R be such that $\text{tbl}(\text{diff}[M_L, M_R]) \in \mathcal{S}$, and $\text{diff}[M'_L, M'_R]$ be such that $\text{tbl}(\text{diff}[M'_L, M'_R]) \in \mathcal{S}$ with $\text{snd}(D)\{x^R \mapsto M'_R\} \Downarrow =_{\text{E}} \text{true}$. Note that the first element exists since we know that $\text{tbl}(M_L) \in \mathcal{S}_{\text{fst}}$. Regarding the second one, we know that it exists relying on our hypothesis $K \Downarrow \uparrow$. Therefore, we consider the two following bi-configurations:

- $K'_0 = (\{i : B_{\text{then}}\{x^L \mapsto M_L, x^R \mapsto M'_R\}\} \cup \mathcal{P}; \Phi; \mathcal{S}; i);$ and
- $K' = (\{i : \mathcal{T}_1(B_{\text{then}})\{x^L \mapsto M_L, x^R \mapsto M'_R\}\} \cup \mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i).$

Since K_0 is separated, we have that K'_0 is separated too, and we have that $\text{fst}(K'_0) = K'_L$. Relying on the fact that our transformation is compatible with substitution, we deduce that $\mathcal{T}_1(K'_0) = K'$. Then, it is easy to see that $K \xrightarrow{\tau\tau} K'$, and thus $K' \Downarrow \uparrow$.

The case of the rule GET-E can be done in a similar way. This allows us to conclude. $\square$

Lemma 3: Let B_0 be a separated closed bi-process such that $\mathcal{T}_1(B_0) \Downarrow \uparrow$, and $T \in \overline{\text{traces}}(\mathcal{T}_1(B_0))$. There exists $T_{\text{fst}} \in \text{traces}(\text{fst}(B_0))$ such that $T_{\text{fst}} \cong \text{fst}(T)$. A similar result holds for snd .

Proof. We prove the result for fst , starting with the following claim.

Claim. Let K_0 be a separated bi-configuration, $K = \mathcal{T}_1(K_0)$ such that $K \Downarrow \uparrow$. If $K \xrightarrow{\alpha} K'$ for some α and some K' with a rule different from GET-T and GET-E, then there exists K'_0 a separated bi-configuration such that $K' = \mathcal{T}_1(K'_0)$, and $\text{fst}(K_0) \xrightarrow{\alpha} \text{fst}(K'_0)$. A similar result holds when considering $K \xrightarrow{\tau\tau} K'$ corresponding to the execution of two `get` instructions in a row. In such a case, we have that $\text{fst}(K_0) \xrightarrow{\tau} \text{fst}(K'_0)$.

Note that the result stated above allows us to prove the statement in the lemma by induction on the length of the execution trace T . Now, it remains to establish this claim. We do a case analysis on the rule(s) used to perform $K \rightarrow K'$.

Case IN. In such a case, we have that:

- $K_0 = (\{i : \text{in}(c, \text{diff}[x^L, x^R]); P\} \cup \mathcal{P}; \Phi; \mathcal{S}; i);$
- $K = (\{i : \text{in}(c, \text{diff}[x^L, x^R]); \mathcal{T}_1(P)\} \cup \mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i);$
- $\alpha = \text{in}(c, R)$ for some recipe R such that $R\Phi_{\text{fst}} \Downarrow =_{\text{E}} M$, and $R\Phi_{\text{snd}} \Downarrow =_{\text{E}} M'$;
- $K' = (\{i : \mathcal{T}_1(P)\{x^L \mapsto M, x^R \mapsto M'\}\} \cup \mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i).$

We consider the following bi-configuration:

$$K'_0 = (\{i : P\{x^L \mapsto M, x^R \mapsto M'\}\} \cup \mathcal{P}; \Phi; \mathcal{S}; i)$$

Since K_0 is separated, we have that K'_0 is separated too. Since $K \xrightarrow{\alpha} K'$ and $K \Downarrow \uparrow$, we have that $K' \Downarrow \uparrow$. Relying on the fact that our transformation is compatible with substitution, we deduce that $K' = \mathcal{T}_1(K'_0)$. Then, it is easy to see that $\text{fst}(K_0) \xrightarrow{\alpha} \text{fst}(K'_0)$.

Most of the cases can be done in a similar way. We focus now on the most interesting ones, corresponding to the execution of two `get` instructions in a row. We simply consider the most interesting cases, corresponding to the execution of two `get` instructions in a row.

Case GET-T followed by GET-T. In such a case, we have that:

- $K_0 = (\{i : \text{get } \text{tbl}(\text{diff}[x^L, x^R]) \text{ st. } D \text{ in } B_{\text{then}} \text{ else } B_{\text{else}}\} \cup \mathcal{P}; \Phi; \mathcal{S}; i);$
- $K = \mathcal{T}_1(K_0) = (\{i : B'\} \cup \mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i)$ where:

$$B' = \begin{cases} \text{get } \text{tbl}(\text{diff}[x^L, _]) \text{ st. } \text{fst}(D) \text{ in} \\ \quad \text{get } \text{tbl}(\text{diff}[_, x^R]) \text{ st. } \text{snd}(D) \text{ in} \\ \quad \mathcal{T}_1(B_{\text{then}}) \\ \text{else out}(_, \text{diff}[\text{bad}_L, \text{bad}_R]) \\ \text{else} \\ \quad \text{get } \text{tbl}(\text{diff}[_, x^R]) \text{ st. } \text{snd}(D) \text{ in} \\ \quad \text{out}(_, \text{diff}[\text{bad}_L, \text{bad}_R]) \\ \text{else } \mathcal{T}_1(B_{\text{else}}) \end{cases}$$

- $K' = (\{i : \mathcal{T}_1(B_{\text{then}})\{x^L \mapsto M_L, x^R \mapsto M_R\}\} \cup \mathcal{T}_1(\mathcal{P}); \Phi; \mathcal{S}; i).$

We consider the following bi-configuration:

$$K'_0 = (\{i : B_{\text{then}}\{x^L \mapsto M_L, x^R \mapsto M_R\}\} \cup \mathcal{P}; \Phi; \mathcal{S}; i).$$

Since K_0 is separated, we have that K'_0 is separated too. Since $K \xrightarrow{\tau\tau} K'$, and $K \Downarrow \uparrow$, we have that $K' \Downarrow \uparrow$. Relying on the fact that our transformation is compatible with substitution, we deduce that $K' = \mathcal{T}_1(K'_0)$. Then, it is easy to see that $\text{fst}(K_0) \xrightarrow{\tau} \text{fst}(K'_0)$.

Case GET-T followed by GET-E. In such a case, we have that

$$K' = (\{i : \text{out}(c, \text{diff}[\text{bad}_L, \text{bad}_R])\} \cup \mathcal{P}; \Phi; \mathcal{S}; i).$$

It is easy to see that K' diverges, and thus K diverges too, contradicting our hypothesis.

The other cases GET-E followed by GET-E, and GET-E followed by GET-T can be done in a similar way. This allows us to conclude. $\square$

Theorem 2: Let B_0 be a separated closed bi-process such that $\mathcal{T}_1(B_0) \Downarrow \uparrow$. We have that $\text{fst}(B_0) \approx_t \text{snd}(B_0)$.

Proof. Let $T_{\text{fst}} \in \text{traces}(\text{fst}(B_0))$. By Lemma 2, we know that there exists $T \in \overline{\text{traces}}(\mathcal{T}_1(B_0))$ such that $T_{\text{fst}} \cong \text{fst}(T)$. Then, relying on Lemma 3 (considering the case snd), we know that there exists $T_{\text{snd}} \in \text{traces}(\text{snd}(B_0))$ such that $T_{\text{snd}} \cong \text{snd}(T)$. Moreover, since $T \in \overline{\text{traces}}(\mathcal{T}_1(B_0))$ and $\mathcal{T}_1(B_0) \Downarrow \uparrow$, we have that $\text{fst}(T) \cong \text{snd}(T)$. This allows us to conclude that $T_{\text{fst}} \cong \text{fst}(T) \cong \text{snd}(T) \cong T_{\text{snd}}$. Hence, the result. $\square$