Automatic generation of sources lemmas in TAMARIN: towards automatic proofs of security protocols[★]

Véronique Cortier¹, Stéphanie Delaune², and Jannik Dreier¹

¹ Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

² Univ Rennes, CNRS, IRISA, France

Abstract. TAMARIN is a popular tool dedicated to the formal analysis of security protocols. One major strength of the tool is that it offers an interactive mode, allowing to go beyond what push-button tools can typically handle. TAMARIN is for example able to verify complex protocols such as TLS, 5G, or RFID protocols. However, one of its drawback is its lack of automation. For many simple protocols, the user often needs to help TAMARIN by writing specific lemmas, called “sources lemmas”, which requires some knowledge of the internal behaviour of the tool.

In this paper, we propose a technique to *automatically* generate sources lemmas in TAMARIN. We prove formally that our lemmas indeed hold, for arbitrary protocols that make use of cryptographic primitives that can be modelled with a subterm convergent equational theory (modulo associativity and commutativity). We have implemented our approach within TAMARIN. Our experiments show that, in most examples of the literature, we are now able to generate suitable sources lemmas automatically, in replacement of the hand-written lemmas. As a direct application, many simple protocols can now be analysed fully automatically, while they previously required user interaction.

1 Introduction

Security protocols are notoriously subtle to design and analyse. Many different tools have been developed in order to detect flaws and prove security properties such as authentication, secrecy, or privacy. However, even a simple property like secrecy is undecidable in general [9]. Hence several tools focus on the analysis of a decidable fragment, e.g. by bounding the number of sessions (e.g. AVISPA [1], DeepSec [6]). But when considering wider classes of protocols, more general cryptographic primitives, and an unlimited number of sessions, one necessarily goes beyond the decidable fragment, possibly losing termination or even automation.

One popular tool in that direction is ProVerif [4], a push-button tool that has been able to analyse hundred of protocols including e.g. TLS 1.3 [3], the

[★] This work has been partially supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (grant agreement No 714955-POPSTAR and grant agreement No 645865-SPOOC), as well as from the French National Research Agency (ANR) under the project TECAP.

ARINC823 avionic protocol [5], or the Neuchâtel voting protocol [7]. However, ProVerif may fail to prove some protocols because of some internal approximations. In that case, the user must either simplify the model or just give up.

Another approach has been developed in the tool TAMARIN [11]. One key feature of TAMARIN is that it provides an interactive mode: if the tool fails to automatically prove a property by itself, the user may help the tool, for example by writing intermediate lemmas, or by manually guiding the proof search. Thanks to this approach, TAMARIN supports many features that are typically out of reach of many tools (Diffie-Hellman, stateful protocols), and has been able to prove complex protocols such as 5G AKA [2] with exclusive or, group key agreement protocols [13], or Noise framework [10] with Diffie-Hellman keys.

However, the fact that TAMARIN is not fully automatic makes it more difficult to use, at least in the learning phase. In particular, TAMARIN fails to automatically prove some “simple” protocols of the literature such as the well-known Needham-Schroder protocol or the Denning-Sacco protocol. This is a barrier when teaching the tool for example at the university or in summer schools.

Automation in TAMARIN fails in particular if it encounters “partial deconstructions”. To speed up the analysis, TAMARIN computes in advance, for each protocol and intruder fact, all possible origins (called *sources*) of these facts, which are then repeatedly used in later steps of the analysis. However, this pre-computation can stop in an incomplete stage if TAMARIN lacks sufficient information about the origins of some fact(s). In practice, as soon as TAMARIN encounters such a “partial deconstruction”, it is unlikely that it will be able to prove any interesting property automatically. To solve the issue, the user needs to manually write a “sources lemma” to help TAMARIN. Unfortunately, this manual step has to be done for many protocols, even simple ones.

Our contribution. In this paper, we automate the generation of sources lemmas. The main idea is to provide a systematic analysis of the origins of a term in a protocol. Intuitively, either a term has been forged by the attacker, or it comes from an earlier step in the protocol. To avoid the exploration of too many cases, we base our analysis on “deepest protected” subterms. We prove that the sources lemmas that we generate are indeed true. Our result holds for any protocol provided that the cryptographic primitives can be expressed as a convergent subterm theory (modulo associativity and commutativity) with the finite variant property. This is the case of most standard cryptographic primitives such as symmetric and asymmetric encryptions, as well as signatures.

Interestingly, the correctness of TAMARIN does not rely on the fact that we are able to prove that our sources lemmas hold. TAMARIN will verify them anyway (as done with sources lemmas written by the user). This means that our technique can also be used even in cases where our theoretical justification does not apply. Our theoretical justification simply explains why TAMARIN has a good chance to work. We have implemented our technique in TAMARIN, as a new option `--auto-sources`. With this option, when partial deconstructions are detected, a sources lemma is generated automatically and added to the original

model, so that the user can see it and possibly amend it, if needed. We have validated our approach with two kind of experiments.

- First, we consider simple protocols of the literature, used as benchmarks for most tools. We modelled a handful of them and ran TAMARIN. Our approach is able to solve all partial deconstructions. Actually, we found out that for these simple examples, this was the only reason they were not entirely automatic, hence thanks to our `--auto-sources` option, TAMARIN can now analyse all these examples automatically.
- We also wanted to evaluate how our technique behaves on more complex protocols and on protocols that have not been specified by ourselves. Hence we considered all the models provided within TAMARIN’s distribution, and that contained “partial deconstructions”. For a large majority of them, our technique successfully close all partial deconstructions and for about a half of them, TAMARIN is now even able to analyse the whole protocol automatically.

Unsurprisingly, complex protocols still require the existing manually written intermediate lemmas. However, our technique considerably improves the degree of automation of TAMARIN, yielding a better trade-off between what can be done automatically, and what needs to be done manually.

2 Overview

We illustrate our technique on a simple challenge-response protocol.

$$\begin{aligned} I &\rightarrow R : \{\text{req}, I, n\}_{\text{pk}(R)} \\ R &\rightarrow I : \{\text{rep}, n\}_{\text{pk}(I)} \end{aligned}$$

The initiator sends a nonce n encrypted with the public key of the responder, and then waits for the corresponding answer, i.e. the nonce n encrypted with his own public key. The symbols `req` and `rep` are constants used to avoid confusion between the two types of messages: they indicate whether the ciphertext corresponds to a request or a reply. In TAMARIN the responder role is as follows:

```
rule Rule_R:
  [ In(aenc{'req', I, x}pk(ltkR)), !Ltk(R, ltkR), !Pk(I, pkI) ]
  --[]-> [ Out(aenc{'rep', x}pkI) ]
```

Intuitively, this rule can be read as follows: at the reception of a message of the form `aenc{'req', I, x}pk(ltkR)`, the agent R (with private key `ltkR`) sends the message `aenc{'rep', x}pkI` on the network to the agent I (with public key `pkI`). Note that there are other rules modelling the Initiator role, as well as the key generation. The latter rule creates the `!Ltk` and `!Pk` facts used here to retrieve the agents’ public and private keys.

This protocol rule models the behaviour of the responder role. It can be triggered arbitrary many times, possibly with different values for x . When loading this model in TAMARIN, it turns out that the proof attempt of e.g. a simple

secrecy property of nonce n does not terminate due to partial deconstructions. In TAMARIN's interactive interface, they are identified by dashed green arrows as shown in Figure 1. The green arrow symbolises a *deconstruction chain*. Deconstruction chains are used in TAMARIN's intruder reasoning to extract values from messages output by the protocol. In this example, TAMARIN tries to extract a fresh value from the message output by the rule **Rule_R** (at the top). TAMARIN has computed that if it can decrypt the output of the rule (rule **d_0_adec**) and then extract the second term (rule **d_0_snd**), it obtains the value $x.7$ (a renaming of the variable x given in the initial rule definition). However, here TAMARIN is unable to continue its deconstruction, as $x.7$ can potentially be any value: directly the desired fresh value, or a pair of values, or an encryption, or something completely different. As this deconstruction is incomplete, it is called a *partial deconstruction*.

Fig. 1. Example of a partial deconstruction

In the above example, TAMARIN does not know anything about the contents of the variable $x.7$, hence, to ensure soundness, it is obliged to consider this case

as a potential source for any value, which leads to an explosion of the number of cases, and often to non termination issues. This is the case here: the rule `Rule_R` producing the `x.7` requires an input, which could itself be the result of (a different instantiation of) the same source, and so on.

To get rid of partial deconstructions, TAMARIN uses *source lemmas*. They are a special type of lemmas which are applied at the precomputation phase. More precisely, after computing the initial *raw sources* without any lemmas, TAMARIN computes the *refined sources* using the source lemmas to hopefully discard partial deconstructions. To ensure that the refined sources are correct, one further has to prove the source lemmas correct, using only the raw sources. This can be done either automatically by TAMARIN or manually in the interactive mode.

The idea behind a source lemma is to provide more information regarding the origin of the message mentioned in the partial deconstruction, i.e., the one corresponding to the variable identified by the dashed green arrow. Going back to our example and assuming that $R(\text{aenc}\{\text{'req'}, I, x\}\text{pk}(\text{ltkR}), x)$ (resp. $I(\text{aenc}\{\text{'req'}, I, n\}\text{pkR})$) is added as a label to the responder rule `Rule_R` (resp. initiator rule), a source lemma could be as follows:

```
lemma typing [sources]:
  "All x m #i. R(m,x)@#i ==> ( (Ex #j. I(m)@#j & #j < #i)
                                | (Ex #j. KU(x)@#j & #j < #i) ) "
```

This lemma says that whenever the responder receives the value `x` inside a message `m` (at time point `#i`), either this message (actually a ciphertext) has been forged by the attacker who therefore knew `x` before, denoted `KU(x)`, or it has been produced (for the first time) by another protocol rule, here the one denoted `I(m)`. Indeed, a quick inspection of the protocol shows that here this is the only option to produce an output having the right format.

When generating the refined sources from the raw sources, TAMARIN applies the source lemmas. In this case, the source lemma above will allow it to learn that `x` is either a nonce (generated by the initiator role) or a message already known by the attacker. This solves the partial deconstruction as the previous source will be refined into two refined sources. The first one is the case where the intruder learns the nonce generated by the initiator, by passing the initiator's message to the responder, and then extracting the nonce like the variable `x.7` above. However, TAMARIN now knows that `x.7` is not any value, but the initiator's nonce. The second case will be discarded by TAMARIN since, if the intruder already knew `x` before, it is useless to extract it again.

3 TAMARIN syntax and semantics

We explain here the syntax and semantics of TAMARIN, as presented in [12, 8], as necessary background for the remainder of the paper.

3.1 Term algebra

Cryptographic messages are represented by a (sorted) term algebra. In TAMARIN, terms are all of sort *msg* and there are two incomparable subsorts *fr* and *pub* used to represent respectively fresh names (e.g. nonces or keys) and public names (e.g. agent names). We assume an infinite set $\mathcal{N}$ of names of each sort and an infinite set $\mathcal{V}$ of variables of each sort as well. A variable x of sort s is denoted $x : s$. The sort *msg* is often omitted, that is, the variable x typically denotes a variable of sort *msg*. Each cryptographic primitive is represented by a function symbol $f : s_1 \times \dots \times s_n \rightarrow s$ that takes n arguments of sort resp. $s_1, \dots, s_n$ and returns a term of sort s . We assume given a *signature* Σ , i.e. a set of function symbols with their arities. Then the set of terms is built from the application of symbols of Σ to names and variables and is denoted $T_\Sigma(\mathcal{N}, \mathcal{V})$. The set of variables occurring in a term t is denoted $\text{vars}(t)$. A term is *ground* if it contains no variable. A substitution θ is *grounding* for t if $t\theta$ is ground.

Example 1. The standard primitives are often expressed by the signature

$$\Sigma_{\text{stand}} = \{\text{enc}(-, -), \text{dec}(-, -), \text{encA}(-, -), \text{decA}(-, -), \text{pk}(-), \langle -, - \rangle, \text{fst}(-), \text{snd}(-)\}$$

where all functions are of sort $\text{msg} \times \dots \times \text{msg} \rightarrow \text{msg}$. They model respectively symmetric encryption and decryption, asymmetric encryption and decryption, and concatenation and (left and right) projections.

The properties of the primitives are reflected through an equational theory E . In TAMARIN, user defined equational theories are given as a convergent rewrite system. TAMARIN additionally supports built-in theories such exclusive or [8] and a set of equations for Diffie-Hellman (DH) exponentiation [12]. The equality modulo associativity and commutativity (AC) is denoted $=_{AC}$ and the normal form of a term t , modulo AC , is denoted $t\downarrow$ (we consider any representative of the normal form of t). Two terms t_1 and t_2 are *unifiable* (modulo AC) if there exists a substitution θ such that $t_1\theta =_{AC} t_2\theta$. *Positions* of a term t are defined as usual considering AC operators as binary symbols. A *subterm* of t is a term t' such that $t' = t|_p$ for some position p .

TAMARIN assumes equational theories that have the finite variant property, that is where all the instances of a given term follow a finite number of different patterns. Formally, a convergent equational theory E has the *finite variant property* if for any term t , there exists a finite number of substitutions $\sigma_1, \dots, \sigma_k$ such that, for any substitution θ , there is $1 \leq i \leq k$, there exists a substitution θ' such that $(t\theta)\downarrow =_{AC} t\sigma_i\theta'$. A particular class of rewriting systems is the class of *subterm* rewriting system. A rewriting system is said subterm if it is defined by a set of equations of the form $l \rightarrow r$ such that r is a subterm of l or a (public) constant. Many cryptographic primitives can be modelled by (convergent) subterm rewriting systems, such as signatures, symmetric and asymmetric encryption, pair, hash, etc. Our theoretical development only consider equational theories that can be defined by a subterm rewriting system, convergent modulo AC , that have the finite variant property. TAMARIN is not limited to subterm equational theories, and actually our approach can be applied in this general setting too relying on Tamarin to establish the correctness of the generated lemmas.

Example 2. Orienting from left to right the equations below yields a subterm convergent rewrite system that is usually used to model concatenation and asymmetric encryption. Here, there is no AC symbol.

$$\text{decA}(\text{encA}(x, \text{pk}(y)), y) = x \quad \text{fst}(\langle x, y \rangle) = x \quad \text{snd}(\langle x, y \rangle) = y$$

In what follows, we will consider sets and multisets. Given a multiset S , $\text{set}(S)$ denotes the set of its elements. The symbol $\subseteq$ denotes the set inclusion. We will write $S \subseteq S'$ even if S and S' are multisets, which is then interpreted as $\text{set}(S) \subseteq \text{set}(S')$. In contrast, $\subseteq^\#$ denotes the multiset inclusion. Similarly, $\cup^\#$ denotes the multiset union and $\setminus^\#$ the multiset difference.

3.2 Transition system

In TAMARIN, a protocol execution is modelled as a transition system where a state contains a multiset of facts, representing the current knowledge of the attacker and the current steps of the protocol, for each agent and each session. Formally, we assume a set of *fact symbols* $\mathcal{F}$ partitioned into *linear* and *persistent* fact symbols. A *fact* is an expression $F(t_1, \dots, t_n)$ where $F \in \mathcal{F}$ and $t_1, \dots, t_n \in T_\Sigma(\mathcal{N}, \mathcal{V})$. Given a multiset of facts F , $\text{lfacts}(F)$ denotes the multiset of its linear facts while $\text{pfacts}(F)$ denotes the multiset of its persistent facts.

Linear facts represent resources that are consumed. TAMARIN includes three pre-defined linear fact symbols: $\text{Fr}(n)$ models the generation of a fresh name n , $\text{Out}(m)$ represents a message m sent over the network by a participant, and $\text{In}(m)$ denotes that the adversary has sent message m , that can then be received by an agent of the protocol. Persistent facts represent facts that remain forever and are not consumed by rules. TAMARIN includes the persistent fact symbol K that models the knowledge of the attacker, as well as $\text{K}^\uparrow$ and $\text{K}^\downarrow$ that allow to distinguish between the terms built by the attacker and those obtained from listening to the network or by decomposing learned messages. Then the protocol may use other user defined facts, that can be either linear or persistent.

The protocol execution is specified through labelled multiset rewriting rules $[l] \multimap [a] \multimap [r]$ where l, a, r are multisets of facts. The multiset l denotes the *premises* of the rule that need to be present in the state in order for the rule to be executed; a denotes the *actions* of the rule (later used to specify properties), while r contains the *conclusions*, added to the state. There are three kinds of rules.

Fresh name generation (FRESH). This is the only rule that can produce facts of the form $\text{Fr}(n)$. Moreover, to ensure freshness, a distinct name n is used for each application.

$$[] \multimap [] \multimap [\text{Fr}(x : fr)]$$

Message deduction rules (MD). They are pre-defined in TAMARIN and represents the attacker's actions.

$$[\text{Out}(x)] \multimap [] \multimap [\text{K}^\downarrow(x)] \quad \text{and} \quad [\text{K}^\uparrow(x)] \multimap [\text{K}(x)] \multimap [\text{In}(x)]$$

model the fact that the attacker can learn any message sent by the protocol and conversely, may send any message of her knowledge. Note that this is the only

rule where the predicate K appears as an action of a rule. The rules

$$\Box \multimap [K^\uparrow(x)] \multimap [K^\uparrow(x : pub)] \quad \text{and} \quad [Fr(x : fr)] \multimap [K^\uparrow(x)] \multimap [K^\uparrow(x : fr)]$$

express respectively that the attacker can learn any public name and can create fresh name on his own. Finally, the attacker can extend his knowledge by applying function symbols. The intuitive rule is:

$$[K(x_1), \dots, K(x_n)] \multimap \Box \multimap [K(f(x_1, \dots, x_n))] \quad \text{for any } f \in \Sigma$$

Actually, this rule is split into two cases in TAMARIN, depending on whether the attacker is building a term, or decomposing it. Formally, for any substitution θ (in normal form), we consider the rule

$$[K^\uparrow(x_1\theta), \dots, K^\uparrow(x_n\theta)] \multimap [K^\uparrow(f(x_1, \dots, x_n)\theta)] \multimap [K^\uparrow(f(x_1, \dots, x_n)\theta)]$$

when $f(x_1, \dots, x_n)\theta$ is in normal form. When the term $f(x_1, \dots, x_n)\theta$ reduces to a subterm of $x_{i_0}\theta$ for some i_0 (remember that we only consider subterm theories), then we consider

$$[K^{\alpha_1}(x_1\theta), \dots, K^{\alpha_n}(x_n\theta)] \multimap [K^\downarrow(f(x_1, \dots, x_n)\theta \downarrow)] \multimap [K^\downarrow(f(x_1, \dots, x_n)\theta \downarrow)]$$

where $\alpha_i = \uparrow$ for all $i \neq i_0$ and $\alpha_{i_0} = \downarrow$. Intuitively, the deduction rule is annotated with $K^\uparrow$ when the attacker applies a “constructor” term such as an encryption and a pair. It can also be annotated with $K^\uparrow$ when the attacker applies a deconstructor (for example, a decryption), if the term cannot be further reduced (for example, the decryption fails). Conversely, the deduction rule is annotated with $K^\downarrow$ when the attacker decomposes a term. Finally, it is possible to switch from $K^\downarrow$ to $K^\uparrow$ thanks to the “coerce” rule:

$$[K^\downarrow(m)] \multimap [K^\uparrow(m)] \multimap [K^\uparrow(m)]$$

for any m in normal form that is not a pair.

Protocol rules. Then the protocol as well as additional attacker capabilities are specified through protocol rules, that are multiset rewriting rules that satisfy some conditions.

Definition 1. A protocol rule is a multiset rewriting rule $[l] \multimap [a] \multimap [r]$ such that

1. it does not contain fresh names and Fr does not occur in r
2. K , $K^\uparrow$, $K^\downarrow$, and Out do not occur in l
3. K , $K^\uparrow$, $K^\downarrow$, In do not occur in r
4. $vars(r) \subseteq vars(l) \cup \{x \in \mathcal{V} \mid x : pub\}$.

The first condition guarantees in particular that fresh names are only produced thanks to the fresh name generation rule. The last three rules are easily met by any rule modelling a protocol step.

Example 3. Going back to our running example, the rule given in Section 2 is a protocol rule where Ltk and Pk are user-defined persistent facts used to model generation of long-term keys. Actually, our model contains the following rule:

$$[Fr(xsk)] \multimap \Box \multimap [!Ltk(xid, xsk), !Pk(xid, pk(xsk)), Out(pk(xsk))]$$

where sk is variable of sort fr , and xid is a variable of sort pub . This protocol rule represents the possibility to generate key pairs $(sk, pk(sk))$ for any identity xid . The public part of the key is revealed to the attacker.

3.3 Execution traces

A set of protocol rules P induces a transition relation $\rightarrow_P$ between states. Namely, we have $S \rightsquigarrow_P^{set(a\theta)} S'$ if there exists a rule $ru \in P \cup MD \cup \{\text{Fresh}\}$ and a grounding substitution θ for ru such that

- $lfacts(l\theta) \subseteq^\# S$, the linear facts of $l\theta$ should be present in S , with enough occurrences,
- $pfacts(l\theta) \subseteq S$,
- and $S' = (S \setminus^\# lfacts(l\theta)) \cup^\# r\theta$. The linear facts of $l\theta$ are removed and all the conclusion facts are added to the state.

Moreover, if the applied rule is the FRESH rule then $r\theta = \{\text{Fr}(n)\}$ and n must be a new name not used earlier. The execution of a protocol is simply modelled by a sequence of transitions. A *trace* of a protocol is the sequence of actions that appear in the execution. Formally, we have that:

$$traces(P) = \{[A_1, \dots, A_n] \mid \emptyset \rightsquigarrow_P^{A_1} \dots \rightsquigarrow_P^{A_n} S'\}.$$

Example 4. Continuing Example 3, the protocol rule modelling key generation can be used twice (or even more) to generate two key pairs for two different identities leading to the following trace:

$$\begin{aligned} \{\} &\rightsquigarrow \{\text{Fr}(ska)\} \rightsquigarrow F_a \cup \{\text{Out}(pk(ska))\} \\ &\rightsquigarrow \{\text{Fr}(skb)\} \rightsquigarrow F_a \cup F_b \cup \{\text{Out}(pk(ska)), \text{Out}(pk(skb))\} \\ &\rightsquigarrow F_a \cup F_b \cup \{K^\downarrow(pk(ska)), \text{Out}(pk(skb))\} \end{aligned}$$

where $F_a = \{!Ltk(A, ska), !Pk(A, pk(ska))\}$, $F_b = \{!Ltk(B, skb), !Pk(B, pk(skb))\}$. Here ska and skb are names of sort fr whereas A, B are public names of sort pub . This corresponds to the application of the FRESH rule followed by the protocol rule to obtain key material for the first agent A and then for a second agent B . The last rule corresponds to an application of an MD rule adding the public key of A to the knowledge of the attacker.

3.4 Properties

Security properties are expressed as properties on the traces of a protocol. TAMARIN offers a first order logic to specify properties. Formulas make use of variables of a novel sort *temp* to reason about when a fact occurs and to be able to express that some event occurs before another one. The full syntax and semantics of the logic is provided in [12]. We provide here only informally the semantics of atomic formulas:

- $F@i$, where i is of sort *temp*, refers to the fact F that occurs in the i^{th} element of the trace;

- $i \doteq j$ expresses that the timepoints i and j are equal;
- $i < j$ expresses that timepoint i occurs before j ;
- $t_1 \approx t_2$ says that t_1 and t_2 are equal (modulo the equational theory).

The first order logic is built from atomic formulas and closed by the boolean connectors $\vee$, $\wedge$, and $\neg$, as well as the quantifiers $\exists$ and $\forall$.

A set of protocol rules P *satisfies* a formula ϕ , denoted $P \models \phi$ if, for any trace $tr \in \text{traces}(P)$, then tr satisfies ϕ .

Example 5. Continuing the running example, a typical lemma expressing nonce secrecy of the challenge is as follows:

lemma nonce_secretcy:

"not(Ex A B s #i #j. (SecretI(A, B, s)@#i & K(s)@#j))"

This requires us to annotate the rule of the Initiator role with the action fact **SecretI**. Then intuitively this lemma expresses that there does not exist any trace such that **SecretI**(A, B, s) occurs at stage i (for some A , B , and s) and the attacker knows s at stage j . If we consider only the three protocol rules mentioned so far (initiator's rule, responder's rule, and key generation), then this security property is satisfied. However, as expected, the same lemma is not satisfied as soon as we model corruption, for example with the following rule.

rule Reveal_ltk: [!Ltk(xid, xsk)] --[RevLtk(xid)]-> [Out(xsk)]

TAMARIN also allows to express diff-equivalence, a refined notion of equivalence. This can be used for example to state that a protocol preserves unlinkability, anonymity, or other privacy properties such as ballot privacy. For example, the fact that Alice remains anonymous is often expressed as the property that $P(\text{Alice}) \sim P(\text{Bob})$. This intuitively says that an adversary should not see the difference when Alice is playing protocol P or Bob is playing protocol P . The formal definition of diff-equivalence can be found in [12]. We do not need to provide it here as our automatically generated lemmas are simple trace properties and do not use diff-equivalence. Note however that our approach applies to protocols with diff-equivalence as well since our generated lemmas also helps TAMARIN to terminate in the case of diff-equivalence properties.

4 Automatically generated sources lemmas

Whenever TAMARIN fails to complete a deconstruction, we aim at providing the tool with a sources lemma that resolves the partial deconstruction. We formalise here our approach and prove it to be correct.

4.1 Definitions

We introduce the notion of *protected* term, which is any term that is headed by a function symbol that is not a pair (because we know the adversary can always open such terms) nor an *AC* symbol (simply because our heuristic does not apply to case of failures due to an *AC* theory).

Definition 2. A protected term t is a term whose head symbol is not $\langle -, - \rangle$ nor an AC symbol. Given a term t and a variable x occurring in t , we say that t' is a deepest protected subterm w.r.t. x if t' is a protected subterm of t that contains x and such that one of the paths from the root of t' to x contains only pair symbols $\langle -, - \rangle$ (except for head symbol at top level).

Intuitively, if t' is a deepest protected subterm w.r.t. x , then the only way to obtain t' is either by extracting it directly from some output, or by building it, in which case x is already known to the attacker.

Example 6. Let $t = \text{enc}(\langle x, \text{enc}(\langle b, x \rangle, k_2) \rangle, k_1)$. There are two deepest protected subterms w.r.t. x , namely t itself and $t' = \text{enc}(\langle b, x \rangle, k_2)$.

We denote by $St_{\text{pair}}(u)$ the set of subterms of u that can be obtained from u simply by projecting. Formally, $St_{\text{pair}}(u)$ is formally defined as

$$St_{\text{pair}}(u) = \begin{cases} \{u\} \cup St_{\text{pair}}(u_1) \cup St_{\text{pair}}(u_2) & \text{if } u = \langle u_1, u_2 \rangle \\ \{u\} & \text{otherwise} \end{cases}$$

Normalised traces. In order to keep track of the origin of a protected subterm, we need to assume that the shape of a term is not modified by the application of the equational theory. Fortunately, since we assume an equational theory with the finite variant property, it is possible to compute in advance the shapes of all the terms obtained after normalisation. Given a set of protocol rules P , TAMARIN computes the variants $\text{Variant}(P)$ of P such that, for any rule $ru \in P$, for any substitution θ , there is $ru' \in \text{Variant}(P)$ and a substitution θ' such that $ru\theta =_E ru'\theta'$ and (ru', θ') is *normalised*, that is, for any fact $F(u')$ occurring in ru' , we have that $(u\theta') \downarrow =_{AC} u'\theta'$. Moreover, $ru' = (ru\sigma) \downarrow$ for some σ .

TAMARIN considers only traces that are *normalised*, i.e. executions of the form $\emptyset \rightsquigarrow_{\text{Variant}(P)}^{A_1} S_1 \cdots \rightsquigarrow_{\text{Variant}(P)}^{A_n} S_n$ and such that:

- the execution involves only rules $ru \in \text{Variant}(P)$ and substitutions θ such that (ru, θ) is normalised;
- pairs are always decomposed before been used, that is, if $K^\uparrow(u)$ appears in the left-hand-side of A_i then $K^\uparrow(t) \in S_{i-1}$ for any $t \in St_{\text{pair}}(u)$ ¹.

We write $P \models_{\text{norm}} \phi$ if for any normalised trace tr of P , tr satisfies ϕ . Then, given a formula ϕ that does not contain the fact $K^\uparrow$ nor $K^\downarrow$, we have $P \models \phi$ if, and only if, $P \models_{\text{norm}} \phi$, which is what is actually checked by TAMARIN. This follows from the soundness of TAMARIN [12].

In some cases, computing the variants $\text{Variant}(ru)$ of a protocol rule ru may introduce new variables on the right of the rule, and thus lead to rules that are not protocol rules (according to Definition 1).

¹ This comes from the fact that, whenever the attacker learns a pair $K^\downarrow(\langle m_1, m_2 \rangle)$, she cannot directly convert it in $K^\uparrow(\langle m_1, m_2 \rangle)$ since the coerce rule does not apply to terms headed with a pair. Hence it is necessary to decompose it first (with $K^\downarrow$ rules) and then reconstruct it (with $K^\uparrow$ rules).

Example 7. The rule $[\text{In}(\text{decA}(x, y))] \multimap [\text{Out}(x)]$ is a protocol rule. However, one of its variant is $[\text{In}(z)] \multimap [\text{Out}(\text{encA}(z, \text{pk}(y)))]$ which is not a protocol rule according to Definition 1.

However, such cases correspond to badly defined protocols and TAMARIN typically raises a warning in this case. Hence, in what follows, we consider *well-formed protocol rules* P , that is such that $\text{Variant}(P)$ is still a set of protocol rules. In practice, protocol rules representing a protocol are indeed well-formed.

4.2 Algorithm

Given a set P of protocol rules, TAMARIN first computes its variants $\text{Variant}(P)$. It then precomputes sources as already explained. Whenever TAMARIN fails to complete a deconstruction, it returns the partial deconstruction. For the moment, assume that from there we can extract a rule $ru = [l] \multimap [a] \multimap [r]$ of $\text{Variant}(P)$ and a variable x for which the deconstruction has failed (in practice there might be multiple composed rules, as explained below, but the approach is similar). It must be the case that x appears in some fact of l .

For each fact symbol F occurring in P , for each rule ru of $\text{Variant}(P)$, and each (deepest) protected subterm t occurring in of ru , we assume new fact symbols $\text{Left}_{F, ru, t}$ and $\text{Right}_{F, ru, t}$ that will be used to further annotate the rules of $\text{Variant}(P)$. These facts will appear only in the sources lemmas we generate.

The sources lemma $\text{SourceLemma}(P, ru, x)$ associated to a failed deconstruction on variable x and rule ru for protocol P is defined by Algorithm 1. Intuitively, we first look for any occurrence of x in the premisses of ru , under a (deepest) protected term t_1 and we annotate the rule ru with $\text{Left}_{F, ru, t_1}(t_1, x)$. Then we look for all facts in the conclusions of a rule ru' that may have produced t_1 , that is that contain a term t_2 that can be unified with t_1 and we annotate ru' with $\text{Right}_{F', ru', t_1}(t_2)$. Finally, we generate the formula that says that if we have $\text{Left}_{F, ru, t_1}(y, x)$ at some step i , then either x is already known to the attacker, that is $\text{K}(x)$ holds at an earlier step, or y has been obtained from the protocol, that is $\text{Right}_{F', ru', t_1}(y)$ holds at some earlier step.

We can show that under our assumptions the generated sources lemmas always hold, which explains why TAMARIN is usually able to prove them.

Theorem 1. *Given a set of well-formed protocol rules P , a rule $ru \in \text{Variant}(P)$, a variable x occurring in ru , and ϕ returned by $\text{SourceLemma}(\text{Variant}(P), ru, x)$, then ϕ is satisfied by $\text{Variant}(P)$, that is $\text{Variant}(P) \models_{\text{norm}} \phi$.*

4.3 Dealing with composed rules

Actually, during the precomputations, TAMARIN might compute the composition of several rules. For example, when a rule ru_1 depends on a rule ru_2 in the sense that ru_1 can only be executed if ru_2 has been executed previously, TAMARIN will return the composition of both, not only ru_1 . This yields bigger steps and it allows TAMARIN to prove lemmas more quickly.

Algorithm 1 SourceLemma(P, ru, x)

Input: $P, ru = [l] \multimap [a] \multimap [r], x$
for all t_1 deepest protected term w.r.t. x that is subterm of $F(v) \in l$ **do**
 % we annotate ru with the fact that x may provide from t_1
 $a := a \cup \{\text{Left}_{F, ru, t_1}(t_1, x)\}$
 % then we identify from which facts t_1 may provide.
for all rule $ru' = [l'] \multimap [a'] \multimap [r'] \in P$ **do**
if t_1 unifiable with t_2 modulo AC for some t_2 protected subterm in $F'(v') \in r'$
then
 % we annotate ru' with the fact that t_2 may be used to produce x
 $a' := a' \cup \{\text{Right}_{F', ru', t_1}(t_2)\}$
end if
end for
 Let ϕ the formula defined as follows

$$\forall y, x, i \text{ Left}_{F, ru, t_1}(y, x) @ i \implies \begin{aligned} & (\exists k \text{ Right}_{F', ru'_1, t_1}(y) @ k \wedge k \leq i) \\ & \vee \dots \\ & \vee (\exists k \text{ Right}_{F', ru'_n, t_1}(y) @ k \wedge k \leq i) \\ & \vee (\exists k K^\dagger(x) @ k \wedge k \leq i) \end{aligned}$$

return ϕ
end for

Thus, the sources computed by TAMARIN are actually composed variants of initial protocol rules. Formally, given two rules $ru_1 = [l_1] \multimap [a_1] \multimap [r_1]$ and $ru_2 = [l_2] \multimap [a_2] \multimap [r_2]$, we define the *composition* of ru_1 and ru_2 w.r.t. θ , denoted $ru_1 \circ_\theta ru_2$ as the rule $[l] \multimap [a] \multimap [r]$ defined as follows:

$$l = l_1 \theta \cup^\# (l_2 \theta \setminus^\# r_1 \theta), \quad a = a_1 \theta \cup a_2 \theta, \quad \text{and} \quad r = (r_1 \theta \setminus^\# l_2 \theta) \cup^\# r_2 \theta.$$

We denote $ru_1 \circ_\theta ru_2 \circ_\theta \dots \circ_\theta ru_k$ the rule ru obtained by iterating $k - 1$ compositions: $ru = ((ru_1 \circ_\theta ru_2) \circ_\theta \dots) \circ_\theta ru_k$. Since the rules do not share any variable, θ is just the union of substitutions θ_i where the domain of θ_i is the set of variables of ru_i . It is easy to check that compositions of protocol rules yield protocol rules. Not all compositions are computed by TAMARIN, but we do not need to characterise which compositions are considered exactly. We simply show that any sources lemma generated from a composed rule is also sound.

Algorithm 2 SourceLemmaComp(P, ru, x)

Input: $P, ru = ru_1 \circ_\theta ru_2 \circ_\theta \dots \circ_\theta ru_k, x$
let l, a, r such that $ru = [l] \multimap [a] \multimap [r]$
for all position p such that there exists $F(v) \in l$ such that $v|_p = x$ **do**
for all i such that $F(v) = F(v_i \theta)$ with $F(v_i)$ in the premisses of ru_i **do**
if p is a position of v_i **then**
 call SourceLemma($P, ru_i, v_i|_p$)
end if
end for
end for

Algorithm 2 describes how to generate a sources lemma from a composed rule. The idea is simply to identify, given a variable x , for which the partial deconstruction is incomplete, at which positions x appears in the composed rule ru . Then whenever the position exists in the some rule ru_i used for composition, we generate the sources lemmas based on this rule. Algorithm 2 is well defined only if whenever $\text{SourceLemma}(P, ru_i, v_i|_p)$ is called, then $v_i|_p$ is a variable. This follows from the fact that $v_i\theta|_p = x$ is a variable (with the notations of Algorithm 2).

Theorem 2. *Given a set of well-formed protocol rules P , a composed rule $ru = ru_1 \circ_\theta ru_2 \circ_\theta \dots \circ_\theta ru_k$ with $ru_i \in \text{Variant}(P)$, a variable x occurring in ru , and ϕ returned by $\text{SourceLemmaComp}(\text{Variant}(P), ru, x)$, then $\text{Variant}(P) \models_{\text{norm}} \phi$.*

5 Implementation and experimental evaluation

We have implemented our approach in TAMARIN version 1.6.0 [15]. The automatic generation of source lemmas is activated using the command line option `--auto-sources`. When TAMARIN is called with this option, it will first load the theory and run the pre-computations normally (in particular compute rule variants and sources). If TAMARIN is called using `--auto-sources`, and the theory does not contain a sources lemma but has partial deconstructions, our new algorithm is executed on the computed rule variants to generate a new sources lemma, which is then added to the theory, as well as the required rule annotations. In the interactive mode, the user can inspect the generated lemma and annotations, and prove lemmas as usual. He can also download the modified theory if he wants to export the lemma, or modify it. In the automatic mode, TAMARIN directly tries to prove the generated sources lemma. When showing the results, TAMARIN displays the sources lemma among the other lemmas, and whether it managed to prove it.

Heuristic. Our first experiments using Algorithm 2 showed that, for some examples, the generated lemmas, while true, caused TAMARIN to loop in the precomputations. This happened when the algorithm considered the case where a fact in the premises of a rule might have been produced by a fact in the conclusion of the same rule. Hence, we have implemented an additional check that ignores this case, should it arise. This means that the generated lemmas could potentially be false, however we did not observe this in practice. In particular, the examples that looped can now be proven correct. Note that this does not contradict our theorems, as our lemmas are not minimal - we consider potentially too many cases, so removing some (unnecessary) ones can still result in a correct lemma.

Evaluation. To evaluate the effectiveness of our approach, we selected several classical examples from the SPORE library of cryptographic protocols [14] and checked for standard properties such as secrecy of the exchanged key and mutual (injective and non-injective) authentication. Because of partial deconstructions, many of them were not entirely automatically verifiable in TAMARIN previously (except for extremely simple examples such as CCITT with only one message).

Protocol Name	Partial Dec.	Resolved	Automatic	Time
Andrew Secure RPC	14	✓	✓	42.8s
Modified Andrew Secure RPC	21	✓	✓	134.3s
BAN Concrete Andrew Secure RPC	0	-	✓	10.6s
Lowe modified BAN Andrew Secure RPC	0	-	✓	29.8s
CCITT 1	0	-	✓	0.8s
CCITT 1c	0	-	✓	1.2s
CCITT 3	0	-	✓	186.1s
CCITT 3 BAN	0	-	✓	3.7s
Denning Sacco Secret Key	5	✓	✓	0.8s
Denning Sacco Secret Key - Lowe	6	✓	✓	2.7s
Needham Schroeder Secret Key	14	✓	✓	3.6s
Amended Needham Schroeder Secret Key	21	✓	✓	7.1s
Otway Rees	10	✓	✓	7.7s
SpliceAS	10	✓	✓	5.9s
SpliceAS 2	10	✓	✓	7.3s
SpliceAS 3	10	✓	✓	8.7s
Wide Mouthed Frog	5	✓	✓	0.6s
Wide Mouthed Frog Lowe	14	✓	✓	3.5s
WooLam Pi f	5	✓	✓	0.6s
Yahalom	15	✓	✓	3.1s
Yahalom - BAN	5	✓	✓	0.9s
Yahalom - Lowe	21	✓	✓	2.2s

Table 1. SPORE examples. “Partial Dec.” indicates the number of partial deconstructions, “Resolved” indicates whether our auto-generated lemmas resolve them, and can be proven correct by TAMARIN. “Automatic” means that our auto-generated lemmas are then sufficient to directly prove or disprove the desired security properties.

The results are presented in Table 1, the TAMARIN models are available in the directory `examples/features/auto-sources/spore` of the TAMARIN repository [15]. Our approach succeeded in all cases.

To see whether our approach works on more complicated examples, we selected all files from the TAMARIN github repository [15] that contained lemmas annotated with `sources`, and that were not marked as “experimental” or “work in progress”. It turned out that in some cases these examples did not actually contain any partial deconstructions, and that these “sources” lemmas were actually used to prove other protocol invariants. As our approach is only meant to handle partial deconstructions, we removed these examples from the set. Table 2 summarises our results on the remaining examples, the files can be found in the directory `examples/features/auto-sources/tamarin-repo` of the TAMARIN repository [15].

It turns out that our algorithm still succeeds in generating successful sources lemmas in the majority of cases, in the sense that the sources lemma resolve all the partial deconstructions and can be proved by TAMARIN. Our examples

Name	Partial Dec.	Resolved	Automatic	Time (new)	Time (previous)
Feldhofer (Equivalence)	5	✓	✓	3.8s	3.5s
NSLPK3	12	✓	✓	1.8s	1.8s
NSLPK3 untagged	12	✓	✗ ¹	-	-
NSPK3	12	✓	✓	2.4s	2.2s
JCS12 Typing Example	7	✓	✗ ²	0.3s	0.2s
Minimal Typing Example	6	✓	✓	0.1s	0.1s
Simple RFID Protocol	24	✓	✗ ²	0.7s	0.5s
StatVerif Security Device	12	✓	✓	0.3s	0.4s
Envelope Protocol	9	✓	✗ ²	25.7s	25.3s
TPM Exclusive Secrets	9	✓	✗ ²	1.8s	1.8s
NSL untagged (SAPIC)	18	✓	✓	4.3s	19.9s
StatVerif Left-Right (SAPIC)	18	✓	✓	28.8s	29.6s
TPM Envelope (Equivalence)	9	✗ ³	-	-	-
5G AKA	240	✗	-	-	-
Alethea	30	✗	-	-	-
PKCS11-templates	68	✗	-	-	-
NSLPK3XOR	24	✗	-	-	-
Chaum Offline Anonymity	128	✗	-	-	-
FOO Eligibility	70	✗	-	-	-
Okamoto Eligibility	66	✗	-	-	-

Table 2. Examples from TAMARIN repository. ¹ The sources lemma needs to be annotated with **reuse** for the following lemmas to be proven automatically. ² The file contains further intermediate lemmas annotated with reuse. ³ The generated lemma removes all partial deconstructions, however TAMARIN does not terminate while trying to prove its correctness automatically.

include protocols with equivalence properties and SAPIC-generated² theories. However, as the examples are more complex, even with a correct sources lemma, TAMARIN does not always succeed in proving all other lemmas fully automatically.

We also analysed the examples where our algorithm failed to generate a correct sources lemma. The reasons turned out to be a too complex equational theory (e.g., FOO and Okamoto, using blind signatures, or NSLPK3XOR and Chaum using XOR), or a complex protocol model where the partial deconstructions stem from the handling of state facts, which escapes our definition of protected subterms (5G AKA, Alethea, PKCS'11). We only encountered one example where the algorithm generated a lemma resolving the partial deconstructions, but TAMARIN was unable to (automatically) verify its correctness.

When our approach succeeds, the verification times are close to timings measured using the manual sources lemmas. All timings have been measured on a standard laptop (Core i7, 16GB RAM, Ubuntu 18.04).

² SAPIC translates from applied pi models to TAMARIN theories.

6 Conclusion

We have provided a technique that allows to automatically generate sources lemmas in TAMARIN, which otherwise had to be written by the user. In return, most simple protocols can now be analysed automatically with TAMARIN.

As future work, we plan to look for even more automation. First, in several cases where our sources lemmas solve the partial deconstructions but are not yet sufficient to prove the security properties specified by the user, we are actually close to full automation. What is missing is simply to indicate to TAMARIN that it should reuse one of the properties (e.g. secrecy of some long-term key) to prove another property (e.g. authentication). We plan to investigate how to automate these “re-use” annotations, without increasing the complexity of the tool.

Our result holds for subterm convergent theories (modulo AC) that have the variant property. However, our algorithm does not generate lemmas for terms headed with an AC symbol (for example exclusive or) as the resulting lemmas would be false in most cases. Hence, manual sources lemmas are still necessary. We plan to explore how to extend our result to tackle this case, which may require to write more complex sources lemmas, e.g. to account for all possible decompositions induced by the exclusive or operator.

Our algorithm also fails when the model uses state facts in such a way that the variables in question do not occur within protected subterms. By generalising the notion of protected subterms, we hope to also cover these cases.

Thanks to our sources lemma, the automation of TAMARIN has improved, in particular on simple protocols. It would be interesting to compare extensively the tools ProVerif and TAMARIN, in order to identify on which cases they are both automatic, and on which kind of protocols, one of the two tools is more likely to conclude automatically. This should also provide directions to improve the automation of both tools.

References

1. A. Armando, D. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuellar, P. Hanks, D. Drielsma, P.-C. Héam, O. Kouchnarenko, J. Mantovani, S. Mödersheim, D. von Oheimb, M. Rusinowitch, J. Santiago, M. Turuani, L. Viganò, and L. Vigneron. The AVISPA Tool for the automated validation of internet security protocols and applications. In K. Etessami and S. Rajamani, editors, *17th International Conference on Computer Aided Verification, CAV'2005*, volume 3576 of *Lecture Notes in Computer Science*, pages 281–285, Edinburgh, Scotland, 2005. Springer.
2. D. Basin, J. Dreier, L. Hirschi, S. Radomirovic, R. Sasse, and V. Stettler. A formal analysis of 5g authentication. In *25th ACM Conference on Computer and Communications Security (CCS'18)*, 2018.
3. K. Bhargavan, B. Blanchet, and N. Kobeissi. Verified models and reference implementations for the tls 1.3 standard candidate. In *IEEE Symposium on Security and Privacy (S&P'17)*, pages 483–503, San Jose, CA, 2017.
4. B. Blanchet. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *14th IEEE Computer Security Foundations Workshop (CSFW-14)*, pages 82–96, Cape Breton, Nova Scotia, Canada, June 2001. IEEE Computer Society.

5. B. Blanchet. Symbolic and computational mechanized verification of the arinc823 avionic protocols. In *30th IEEE Computer Security Foundations Symposium (CSF'17)*, pages 68–82, Santa Barbara, CA, USA, 2017.
6. V. Cheval, S. Kremer, and I. Rakotonirina. Deepsec: Deciding equivalence properties in security protocols - theory and practice. In *Proceedings of the 39th IEEE Symposium on Security and Privacy (S&P'18)*, pages 525–542. IEEE Computer Society Press, May 2018.
7. V. Cortier, D. Galindo, and M. Turuani. A formal analysis of the neuchâtel e-voting protocol. In *3rd IEEE European Symposium on Security and Privacy (EuroSP'18)*, pages 430–442, London, UK, April 2018.
8. J. Dreier, L. Hirschi, S. Radomirovic, and R. Sasse. Automated Unbounded Verification of Stateful Cryptographic Protocols with Exclusive OR. In *CSF 2018*, pages 359–373, 2018.
9. N. Durgin, P. Lincoln, J. Mitchell, and A. Scedrov. Undecidability of bounded security protocols. In *Workshop on Formal Methods and Security Protocols*, Trento, Italia, 1999.
10. G. Girol, L. Hirschi, R. Sasse, D. Jackson, C. Cremers, and D. Basin. A spectral analysis of Noise: A comprehensive, automated, formal analysis of Diffie-Hellman protocols. In *Usenix Security*, 2020.
11. S. Meier, B. Schmidt, C. Cremers, and D. Basin. The TAMARIN Prover for the Symbolic Analysis of Security Protocols. In *Computer Aided Verification, 25th International Conference, CAV 2013, Princeton, USA, Proc.*, volume 8044 of *LNCS*, pages 696–701. Springer, 2013.
12. B. Schmidt, S. Meier, C. J. F. Cremers, and D. A. Basin. Automated Analysis of Diffie-Hellman Protocols and Advanced Security Properties. In *CSF 2012*, pages 78–94, 2012.
13. B. Schmidt, R. Sasse, C. Cremers, and D. Basin. Automated verification of group key agreement protocols. In *IEEE Symposium on Security and Privacy (S&P'14)*, 2014.
14. Security protocols open repository. <http://www.lsv.fr/Software/spore/>. Accessed on 04/24/2020.
15. Main source code repository of the tamarin prover for security protocol verification. <https://github.com/tamarin-prover/tamarin-prover>. Accessed on 12/06/2019.

A Proofs of Theorems 1 and 2

Theorem 1. *Given a set of well-formed protocol rules P , a rule $ru \in \text{Variant}(P)$, a variable x occurring in ru , and ϕ returned by $\text{SourceLemma}(\text{Variant}(P), ru, x)$, then ϕ is satisfied by $\text{Variant}(P)$, that is $\text{Variant}(P) \models_{\text{norm}} \phi$.*

Proof. Let P be a set of protocol rules, $ru \in \text{Variant}(P)$ and a variable x occurring in ru , let ϕ be a formula returned by $\text{SourceLemma}(\text{Variant}(P), ru, x)$. The rule ru is of the form $[l] \multimap [a] \multimap [r]$ and ϕ is of the form:

$$\forall y, x, i \text{ Left}_{F, ru, t_1}(y, x) @ i \implies \begin{array}{l} (\exists k \text{ Right}_{F', ru'_1, t_1}(y) @ k \wedge k < i) \\ \vee \dots \\ \vee (\exists k \text{ Right}_{F', ru'_n, t_1}(y) @ k \wedge k < i) \\ \vee (\exists k \text{ K}^\uparrow(x) @ k \wedge k < i) \end{array}$$

for some t_1 deepest protected term w.r.t. x , subterm of $F(t) \in l$. By definition of a deepest protected subterm, $t_1|_p = x$ for some position p and there are only pairs along the path p (except at position ϵ).

Let tr be a normalised trace of $\text{Variant}(P)$. Let us show that tr satisfies ϕ .

$$tr = \emptyset \rightsquigarrow^{A_1} S_1 \dots \rightsquigarrow^{A_{n-1}} S_{n-1} \rightsquigarrow^{A_n} S_n$$

Let i be such that $\text{Left}_{F,ru,t_1}(m,n) \in S_i$ for some terms m,n . Then the i^{th} applied rule must be the rule ru in $\text{Variant}(P)$ mentioned above which has the form:

$$ru = \{[F(t)] \cup l'\} - [\text{Left}_{F,ru,t_1}(t_1, x) \cup a'] \rightarrow [r]$$

Moreover, there exists a substitution σ_i in normal form (the one used to instantiate ru) such that $m =_{AC} (t_1\sigma_i)\downarrow$ and $n =_{AC} x\sigma_i\downarrow$. Since the trace is normalised, $m =_{AC} t_1\sigma_i$ and $n =_{AC} x\sigma_i$. Let $u =_{AC} (t\sigma_i)\downarrow$. Again, we have $u =_{AC} t\sigma_i$. Since t_1 is a subterm of t and t_1 is not headed by an AC symbol, we have that m is a subterm of u (modulo AC). Moreover $F(u) \in S_{i-1}$ by definition of the application of a rule.

Let $j < i$ be the first occurrence of j such that m (modulo AC) is a subterm of a fact in S_j and consider the j^{th} rule that has been applied.

- Either this rule is a rule ru'' in $\text{Variant}(P)$ of the form

$$ru'' = [l''] - [a''] \rightarrow [\{F'(w)\} \cup r'']$$

and there exists σ_j in normal form (the substitution used to instantiate ru'') such that m (modulo AC) is a subterm of $u' = (w\sigma_j)\downarrow$. Since the trace is normalised, $(w\sigma_j)\downarrow =_{AC} w\sigma_j$. Let p' be the position at which m occurs in $w\sigma_j$, i.e. such that $w\sigma_j|_{p'} =_{AC} m$.

- Either p' is a path of w that does not end on a variable. Then $w|_{p'} = w'$ with w' a protected subterm of w .
We have that $w'\sigma_j =_{AC} m =_{AC} t_1\sigma_i$ thus w' and t_1 are unifiable (modulo AC) thus we have annotated ru'' , that is, $\text{Right}_{F',ru'',t_1}(w') \in a''$, which concludes this case.
- Or p' is a path of w that ends on a variable or is not a path at all. Then there must exist a variable y in w such that m (modulo AC) is a subterm of $y\sigma_j$. Then y also appears in some premise fact $F''(w'')$, thanks to the definition of a protocol rule and the fact that the variant rules are still protocol rules. Therefore m (modulo AC) is a subterm of a fact in S_{j-1} (since $(w''\sigma_j)\downarrow =_{AC} w''\sigma_j$), which contradicts the minimality of j .
- Or the rule is one of the MD rules. Since m is a protected term, the rule cannot be $\square - [K^\uparrow(x)] \rightarrow [K^\uparrow(x : pub)]$ nor $[Fr(x : fr)] - [K^\uparrow(x)] \rightarrow [K^\uparrow(x : fr)]$ since these two rules only generate names. By minimality of j , it cannot be the rule $[Out(x)] - \square \rightarrow [K^\downarrow(x)]$, nor $[K^\uparrow(x)] - [K(x)] \rightarrow [In(x)]$, nor the rule $[K^\downarrow(x)] - [K^\uparrow(x)] \rightarrow [K^\uparrow(x)]$ either. So it must be the deduction rule, either in the $K^\uparrow$ version or in the $K^\downarrow$ version.

- Either it is the rule

$$[K^\uparrow(x_1\theta), \dots, K^\uparrow(x_n\theta)] \multimap [K^\uparrow(f(x_1, \dots, x_n)\theta) \multimap [K^\uparrow(f(x_1, \dots, x_n)\theta)]]$$

with $f(x_1, \dots, x_n)\theta$ in normal form. We have $K^\uparrow(x_1\theta), \dots, K^\uparrow(x_k\theta) \in S_{j-1}$. Then, by minimality of j , and since m is not headed with an AC symbol, we must have $m =_{AC} t_1\sigma_i =_{AC} f(x_1\theta, \dots, x_k\theta)$, otherwise we would have that m is subterm of some $x_i\theta$ hence subterm of S_{j-1} or m is a constant, which cannot be the case since m is a protected subterm. Remember that $x\sigma_i$ is a subterm at position $p = i_0.p'$ (for some i_0) of t_1 such that there are only pairs along p' , that is, $x\sigma_i \in St_{\text{pair}}(x_{i_0}\theta)$. Since the trace is normalised (i.e. pairs are decomposed before being used), we get that $K^\uparrow(x\sigma_i) \in S_{j-1}$, that is $K^\uparrow(n) \in S_{j-1}$. Now, by inspection of the rules, we notice that the only way to obtain $K^\uparrow(t)$ in a state is through a rule annotated by $K^\uparrow(t)$, hence we can conclude that $K^\uparrow(n)$ appears in one of the actions of an earlier rule.

- Or the rule

$$[K^{\alpha_1}(x_1\theta), \dots, K^{\alpha_n}(x_n\theta)] \multimap [K^\downarrow(f(x_1, \dots, x_n)\theta \downarrow) \multimap [K^\downarrow(f(x_1, \dots, x_n)\theta \downarrow)]]$$

has been applied, with $f(x_1, \dots, x_k)\theta$ that can be reduced at top level. Since the equational theory is a subterm theory, it must be the case that $m = (f(x_1, \dots, x_k)\theta) \downarrow$ is a subterm of one of the $x_i\sigma$, hence m is a subterm of a fact of S_{j-1} , which contradicts the minimality of j . $\square$

Theorem 2. *Given a set of well-formed protocol rules P , a composed rule $ru = ru_1 \circ_\theta ru_2 \circ_\theta \dots \circ_\theta ru_k$ with $ru_i \in \text{Variant}(P)$, a variable x occurring in ru , and ϕ returned by $\text{SourceLemmaComp}(\text{Variant}(P), ru, x)$, then $\text{Variant}(P) \models_{\text{norm}} \phi$.*

Proof. The correctness of Algorithm 2 is a direct consequence of Theorem 1. Indeed, let ϕ be a formula returned by $\text{SourceLemmaComp}(\text{Variant}(P), ru, x)$. Then ϕ is actually a formula returned by $\text{SourceLemma}(\text{Variant}(P), ru_i, v_i|_p)$ for some $ru_i \in \text{Variant}(P)$ and some variable $v_i|_p$ of ru_i . Applying Theorem 1, we have that $\text{Variant}(P) \models_{\text{norm}} \phi$, hence the conclusion. $\square$