

Automatic verification of transparency protocols (extended version)

Cheval, Vincent; Moreira, José; Ryan, Mark

DOI:

[10.48550/arXiv.2303.04500](https://doi.org/10.48550/arXiv.2303.04500)

License:

Other (please provide link to licence statement)

Document Version

Other version

Citation for published version (Harvard):

Cheval, V, Moreira, J & Ryan, M 2023 'Automatic verification of transparency protocols (extended version)' arXiv. <https://doi.org/10.48550/arXiv.2303.04500>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Automatic verification of transparency protocols (extended version)

Vincent Cheval
INRIA Paris
France
vincent.cheval@inria.fr

José Moreira
Valory AG
Switzerland
jose.moreira.sanchez@valory.xyz

Mark Ryan
University of Birmingham
United Kingdom
m.d.ryan@cs.bham.ac.uk

Abstract—Transparency protocols are protocols whose actions can be publicly monitored by observers (such observers may include regulators, rights advocacy groups, or the general public). The observed actions are typically usages of private keys such as decryptions, and signings. Examples of transparency protocols include certificate transparency, cryptocurrency, transparent decryption, and electronic voting. These protocols usually pose a challenge for automatic verification, because they involve sophisticated data types that have strong properties, such as Merkle trees, that allow compact proofs of data presence and tree extension.

We address this challenge by introducing new features in ProVerif, and a methodology for using them. With our methodology, it is possible to describe the data type quite abstractly, using ProVerif axioms, and prove the correctness of the protocol using those axioms as assumptions. Then, in separate steps, one can define one or more concrete implementations of the data type, and again use ProVerif to show that the implementations satisfy the assumptions that were coded as axioms. This helps make compositional proofs, splitting the proof burden into several manageable pieces. We illustrate the methodology and features by providing the first formal verification of the transparent decryption and certificate transparency protocols with a precise modelling of the Merkle tree data structure.

Index Terms—protocols, transparency protocols, automatic verification, ProVerif, symbolic model

1. Introduction

Many security protocols assume the availability of trusted third parties, such as cloud computing operators. These are entities which are inherently trusted by definition, and whose corruption or misbehavior might be undetected and have catastrophic consequences from a security standpoint. With the advent of web3 and distributed ledger technologies, there is an ever growing interest in the concepts of transparency and accountability of misbehaviour in security protocols: whereas we cannot prevent a malicious action from happening, we can track actions and take the appropriate corrective measures on the entities that have broken the presumed trust assumptions.

The focus of this paper is on transparency. In general terms, we say that a protocol is a *transparency protocol* if (some of) its actions can be publicly monitored by a collection of observers, which can, at a later stage, provide evidence that certain actions occurred. By technical

measures to ensure that the details of such actions are made available to relevant parties (whether the general public, regulators, rights advocacy groups or individuals), the protocol reduces the amount of trust required of trusted third parties, and deters malicious actions.

To motivate our work, we consider two prominent use cases of transparency protocols. On one hand, we consider the Google initiative on *certificate transparency* [1] for monitoring and auditing digital certificate issuance which has become an IETF standard [2]. It enables the detection of illegitimate certificates that have been produced either erroneously or deliberately. On the other hand, we also consider *transparent decryption* [3], a protocol that ensures visibility of decryption requests, and has applications in a variety of areas such as surveillance, data sharing and location-based services, among many others.

Appropriate transparency conditions for a protocol can be defined using an append-only ledger, or a blockchain. We assume that some party is willing to maintain the ledger, and we demonstrate how this can be organised in such a way that the maintainer can produce data demonstrating that it is, indeed, maintaining the log correctly. The ledger contents are intended to be accessible by any party to whom transparency is being offered. Therefore, any such party can determine that the ledger is running correctly (or that it is not running correctly, or that their access to it has been denied).

Automatic verification. Our paper concerns how to verify systems for transparency protocols. That is, how to prove that the system really does guarantee that the actions of monitored parties are made available to observers.

Several tools have been proposed for automated analysis of security protocols. Some of these tools impose restrictions on the protocols in order to achieve termination of the analysis. For example, the tools may assume a bounded number of sessions, like Avispa [4], DeepSec [5], or Akiss [6]. These tools are efficient at finding attacks on small protocols but quickly face state explosion for complex protocols. Hence for large and complex protocols, tools like Tamarin [7] and ProVerif [8] are often preferred. They both offer a flexible framework to model a protocol and its primitives, as well as their security properties. One key feature of Tamarin is that it offers an interactive mode when the tool fails to prove a protocol, while ProVerif typically offers more automation.

We work within the framework of ProVerif. It supports cryptographic primitives including symmetric and asymmetric encryption; digital signatures; hash functions;

bit-commitment; and signature proofs of knowledge. The tool is capable of evaluating secrecy properties, authentication properties, and indistinguishability properties. In ProVerif, protocol analysis is considered with respect to an unbounded number of sessions and an unbounded message space. The tool is capable of attack reconstruction: when a property cannot be proved, an execution trace which falsifies the desired property can often be constructed.

The transparency ledger may be implemented as a Merkle tree, which means that the ledger maintainer can produce proofs that a data item is in the ledger (*proof of presence*), and that the ledger is only being appended to (*proof of extension*). To model this in ProVerif, we use user-defined predicates whose semantics is defined with Horn clauses. ProVerif can work with arbitrary Horn clauses, but adding them often leads to non termination of the ProVerif resolution strategy. Recent work [9] has extended ProVerif with notions of lemmas and axioms, in an effort to address this non-termination issue. However, lemmas and axioms have strong limitations with user-defined, attacker or message predicates. We address these limitations in this paper.

Our contributions. The paper develops new ProVerif capabilities and a methodology for using them. Our contributions are as follows:

- We introduce new capabilities in ProVerif; more precisely, we define semantics and algorithms that allow ProVerif to work with lemmas and axioms that involve user-defined predicates.
- We prove soundness of the algorithms, and implement them in a new version of ProVerif.
- We introduce a new methodology for ProVerif, in which the proof of a protocol can be given based on assumptions about the behaviour of a data type (we code these assumptions as an *interface*); then, in separate steps we formally prove the assumptions hold for one or more concrete realisations of the data type. This methodology would syntactically not be possible without our extension of ProVerif.
- We model two transparency protocols (transparent decryption and certificate transparency), and successfully instantiate the proposed methodology.

The paper is supported by our new version of ProVerif and the ProVerif scripts for transparent decryption, which can be found in [10]. Our code has been reviewed by the ProVerif owners, and it will be incorporated in the next ProVerif release. Detailed definitions and proofs are provided in Appendix.

2. Background and related work

2.1. Transparency protocols

Several recent protocols use a publicly-accessible append-only log data structure to achieve a *transparency property*. One of the earliest and most widely deployed protocol in this category is *certificate transparency* [1]. The core idea of certificate transparency is that certificates are accepted by browsers only if they are accompanied by a proof that they are present in an appropriate log. Insisting on certificates being in a publicly-accessible log means that the existence of the certificate is *transparent*.

It prevents situations in which corrupted CAs issue rogue certificates without being noticed. This idea has been generalised to define more ambitious public-key infrastructures, such as ARPKI [11] and DTKI [12], which aim to make all the infrastructure parties behave transparently.

The log and its proof data. As mentioned, transparency protocols rely on an append-only log. The log is not assumed to be trustworthy; rather, anyone can verify the data it outputs, and if this verification succeeds, then the transparency property is upheld. The *expected behaviour* of the log may vary from protocol to protocol; here, we give a generic example [13]. The log L is organised as an append-only Merkle tree. For our purposes, a Merkle tree storing data R_1, R_2, \dots, R_n is a binary tree whose leaves (when considered in left-to-right order) store the data R_1, R_2, \dots, R_n and whose non-leaf nodes store $H(c_\ell, c_r)$ where H is a hash function and c_ℓ and c_r is the data stored at the left and right child node respectively. The maintainer of L runs three protocols:

- On request, it outputs the current value h stored at L 's root (called the *root tree hash* of L).
- On input R , it outputs data which proves that R is present in L (or it outputs \perp if that is not the case). This data consists of data stored in some of the nodes of L , and its size is $O(\log n)$.
- On input h_1 and h_2 , it outputs data which proves that L previously had the root tree hash h_1 , and subsequently the root tree hash h_2 (or it outputs \perp if this is not the case). This data also consists of data stored in some of the nodes of L , and its size is $O(\log n)$.

The log's behaviour is fully verifiable, and therefore there is no trust assumption on the log maintainer. It can be malicious, and still the security property is upheld.

Verification of transparency protocols. Several previous papers have applied protocol verification techniques to transparency protocols. The papers on ARPKI and DTKI [11], [12] both use the Tamarin prover to prove some security properties. However, both papers make the same huge abstraction: they treat the log as a list. The DTKI paper acknowledges that formalising and modeling the complex data structures of transparency protocols is an unsolved problem. Certificate Transparency was proved using Tamarin in [14] using a simplified setting, for example by modelling the log as a trusted global database shared between agents. Additionally, in that work, proofs of presence and proof of extension were not modeled; instead, they use placeholders in the shared trusted database to act as proofs. In contrast, we model proofs and their data structure precisely.

2.2. Transparent decryption

Transparent decryption is a transparency protocol, aiming to prevent certain decryptions from being performed stealthily; rather, the decryption operation inevitably produces evidence of the fact that the decryption has taken place. This can be used, for example, to support privacy: it can mean that a subject is alerted to the fact that information about them has been decrypted. Among other uses, transparent decryption has been proposed for

accountable execution of search warrants and data interception [3], [15], [16]; data sharing between organisations [17]; in vehicle and IoT data applications [18]. A company has begun building products using these ideas [19]. Transparent decryption is an accountable algorithm in the sense of [20].

In transparent decryption, the decryption key is distributed among a set of agents (called *trustees*); they use their key share only if the required transparency conditions have been satisfied. Typically, the transparency condition can be formulated as the presence of the decryption request in a transparency log [21].

We present a minimal system for transparent decryption below. The system satisfies the basic security property for transparent decryption, which we also detail below.

How it works. More formally, the system works as follows:

- *Subjects* create ciphertexts using a public encryption key ek .
- Shares dk_1, \dots, dk_n of the decryption key are held by *trustees* T_1, \dots, T_n . For example, this might be a threshold decryption system, so that any m out of n trustees are sufficient to decrypt.
- A *decryption requester* G can request the decryption of a ciphertext. This involves recording the request in a log L .
- L is organised as a Merkle tree. This means that the maintainer can issue data that demonstrates it's maintaining L in an append-only fashion (see below).
- *Trustees* are automatic processes which accept ciphertexts and log data as input. The log data attests that certain information has been placed in the log L , and that L has been maintained append-only since it was last seen by the trustee. The trustees verify the log data. If (and only if) the data verifies correctly, a trustee will perform its part in decrypting the relevant ciphertext and output the result.
- Subjects can try to inspect the log contents. If their attempt is successful, they will see from the log whether their ciphertexts have been decrypted or not. If they are not successful (for example, the data is inconsistent or their access is denied) then they should assume that their ciphertexts have been decrypted.

The trust assumption for transparent decryption is that the trustees behave correctly. As mentioned, this means that they perform their part of the decryption if, and only if, the verification of the proofs in the input data is successful.

Security property. We aim to prove the following property:

Suppose an honest subject encrypts a secret s with ek , and later the secret s becomes known by some other party (e.g., any of the mentioned parties, or an attacker, or anyone else). Suppose the subject successfully accesses the log L and successfully verifies the log data. Then the subject sees the decryption request for s in L .

Trustees and their actions. Trustees are designed to be very simple and to have minimal computational requirements, so that their trustworthiness can be established as straightforwardly as possible. They do not have to store

T_i stores: h, dk_i, sk_i	
<ul style="list-style-type: none"> • Input: R, h', π, ρ • Compute: <ul style="list-style-type: none"> – Verify π: R in h' – Verify ρ: h' extends h – result := dec(dk_i, R) – $h := h'$ • Output: result 	<ul style="list-style-type: none"> • Input: v • Compute: <ul style="list-style-type: none"> $r := \text{sign}(sk_i, (v, h))$ • Output: r

Figure 1. Protocols run by trustee T_i . The trustee stores the most recent root tree hash h of the log that it has seen, and a decryption key dk_i and a signing key sk_i share. The protocol on the left inputs a request R and some other parameters, and outputs a decrypted result. The protocol on the right inputs a nonce v , and outputs a signature on (v, h) .

any voluminous data; they store just three data items, and they run two protocols (see Fig. 1).

Trustees can be implemented in a variety of ways. For example, they may be cloud-based software processes run by organisations with a high reputation such as charities and foundations. These organisations can use hardware-based attestation to give further confidence about the binary code trustees are running, and the secure storage and use of their keys. Alternatively, trustees could be implemented on dedicated hardware modules, such as the TPM [22] or Google Titan chip [23], or a RISC-V chip like Open Titan [24], [25].

The system we have described is a minimal one that provides decryption transparency. It could readily be extended to have some additional properties, such as trustee obliviousness (namely, the inability of a trustee to obtain any information about the decryption request or its result), and proper authentication of the decryption request (see, e.g. [15]).

Applications. Transparent decryption can be applied in many areas to enhance privacy. We give some examples.

- Alex can choose to share her location in encrypted form with some nominated friends and family, called *angels* by the app that implements this idea [26]. No-one except her angels can view her location; and she can monitor whether and when they do so.
- Suppose Alex is being investigated by the police. In an effort to establish her innocence, she may choose to hand over her phone. With transparent decryption, Alex can upload her phone contents in encrypted form. Then Alex gets evidence of what part of this uploaded material is decrypted.
- Alex provides *know-your-customer* (KYC) information to her bank, so that if necessary later, it can carry out anti-money laundering procedures. Recent proposals [27] suggest centralising KYC registers, to make the procedure more efficient. With transparent decryption, the ill effects of such centralisation can be mitigated by making money laundering investigations more transparent.

2.3. ProVerif

ProVerif is a software tool for automated reasoning about the security properties of cryptographic protocols.

It was first released in 2002, and has been continuously developed for the last 20 years. It has been used to analyze hundreds of protocols, including major deployed protocols such as TLS [28], Signal [29], Noise [30], avionic protocols [31], and the Neuchâtel voting protocol [32].

A cryptographic protocol in ProVerif is specified as follows:

- Cryptographic primitives (such as symmetric and asymmetric encryption or digital signatures) are specified typically as *reduction rules*, such as this one for public key encryption and decryption: $\text{decr}(k, \text{encr}(\text{enc_key}(k), m)) = m$.
- The behaviour of the protocol participants is described using the process calculus syntax (see below).
- The properties which are to be checked are specified as queries. ProVerif supports different kinds of properties; in this paper, we restrict our attention to reachability and correspondence properties. For example, the correspondence property $\text{event}(ev(x)) \Rightarrow \text{att}(x)$ says that if the event ev occurs with a parameter value x , then the value x was previously known by the attacker.

Syntax. A simplified syntax for the process calculus terms, expressions, events, predicates and processes is displayed in Fig. 2. ProVerif’s calculus also supports additional constructs, e.g. for tables, phases, extended terms, ..., but we omit them for simplicity as our results can be easily generalized to these constructs.

Terms M, N, \dots are built over variables, names and application of *constructor* function symbols from a finite set \mathcal{F}_c . *Destructor* function symbols, from a finite set \mathcal{F}_d , can manipulate terms and must be evaluated in the *assignment* construct. Unlike constructor function symbols, the evaluation of a destructor may *fail*, or in other words, may evaluate to the special constant fail. Typically, in the assignment construct $\text{let } x = D \text{ in } P \text{ else } Q$, the expression D will be evaluated; if its evaluation fails then the process Q will be executed, otherwise the variable will be instantiated by its result and P will be executed. The exact behavior of a destructor function symbol is defined by a list of rewrite rules given by the user (see [9] for the complete definition of the evaluation of an expression). For example,

$$\text{in}(c, y); \text{let } x = \text{decr}(k, y) \text{ in } \text{out}(c, x) \text{ else } \text{out}(c, 0)$$

outputs the plaintext m if a ciphertext of the form $\text{encr}(\text{enc_key}(k), m)$ is given as input, otherwise it outputs 0.

A substitution σ is an assignment of terms to some variables; for example, $\{x \mapsto \text{encr}(k, m)\}$ is a substitution. If M is a term, then $M\sigma$ is the term obtained by replacing any x mapped by the substitution with the term that it maps to. For processes P and facts F , applying the substitution to obtain $P\sigma$ and $F\sigma$ is defined similarly (taking care not to substitute bound variables).

The process calculus also contains standard constructs $\text{out}(N, M); P$ (representing the output of a term M on a channel N), $\text{in}(N, x); P$ (the input on channel N of a message which gets bound to a variable x), $\text{new } a; P$ (the generation of a fresh name a), $P \mid Q$ (the concurrent execution of processes), $\text{event}(ev(M_1, \dots, M_k)); P$ (the

$P, Q ::=$	processes
0	nil
$\text{out}(N, M); P$	output
$\text{in}(N, x); P$	input
$P \mid Q$	parallel composition
$!P$	replication
$\text{new } a; P$	restriction
$\text{let } x = D \text{ in } P \text{ else } Q$	assignment
$\text{let } x_1, \dots, x_n \text{ suchthat } p(M_1, \dots, M_k) \text{ in } P \text{ else } Q$	predicate evaluation
$\text{event}(ev(M_1, \dots, M_k)); P$	event

Figure 2. Syntax of the core language of ProVerif.

recording of event execution), and $!P$ (the concurrent execution of an unbounded number of copies of a process).

Less common is the construct for predicate evaluation, that is, $\text{let } x_1, \dots, x_n \text{ suchthat } p(M_1, \dots, M_k) \text{ in } P \text{ else } Q$. In this construct, the variables x_1, \dots, x_n must occur in the predicate $p(M_1, \dots, M_k)$. If x_1, \dots, x_n can be instantiated, say by a substitution σ , such that $p(M_1, \dots, M_k)\sigma$ holds then $P\sigma$ is executed; otherwise Q is executed. Note that predicate evaluations are most commonly used with a classical if-then-else conditional, corresponding in fact to the evaluation without variables ($n = 0$).

User-defined predicates and clauses. ProVerif allows users to define predicates, and to give their semantics by means of Horn clauses. This is useful for defining predicates on data types. For example, the list data structure can be represented by a constant nil and a constructor cons. One can define a membership predicate mem using the Horn clauses:

$$\begin{aligned} &\forall x, \ell. \text{mem}(x, \text{cons}(x, \ell)) \\ &\forall x, y, \ell. \text{mem}(x, \ell) \rightarrow \text{mem}(x, \text{cons}(y, \ell)). \end{aligned}$$

As variables in Horn clauses are always universally quantified, we will omit writing the quantifier in the rest of this paper.

Definition 1 (derivation). A derivation \mathcal{D} of a fact F from a set of clause \mathcal{C}_{user} is a tree whose nodes are labeled by Horn clauses in \mathcal{C}_{user} and edges are labeled by ground facts such that the incoming edge of the root is labeled by F . For all nodes η in \mathcal{D} labelled by a clause $F_1 \wedge \dots \wedge F_n \wedge \phi \rightarrow C$, there exists a substitution σ such that: (i) the incoming edge of η is labeled by $C\sigma$; (ii) η has n outgoing edges labeled by $F_1\sigma, \dots, F_n\sigma$ respectively; (iii) $\phi\sigma$ is true.

The derivability of facts allows us to define the true statements of a predicate p , denoted $\text{sem}(p)$, as the set of facts $F = p(M_1, \dots, M_n)$ derivable from \mathcal{C}_{user} . Note that only user-defined predicates, equalities and disequalities on terms can occur in the clauses from \mathcal{C}_{user} . As such, the semantics of a user-defined predicate is independent from any protocol.

Example 1. $\text{mem}(a, \text{cons}(b, \text{cons}(a, \text{nil}))) \in \text{sem}(\text{mem})$ as it is derivable by the following derivation with $\sigma_1 = \{x \mapsto a; y \mapsto b; \ell \mapsto \text{cons}(a, \text{nil})\}$ and $\sigma_2 = \{x \mapsto a; \ell \mapsto \text{nil}\}$.

$$\begin{array}{c}
\text{mem}(a, \text{cons}(b, \text{cons}(a, \text{nil}))) \downarrow \\
\boxed{\text{mem}(x, \ell) \rightarrow \text{mem}(x, \text{cons}(y, \ell))} \text{ with } \sigma_1 \\
\text{mem}(a, \text{cons}(a, \text{nil})) \downarrow \\
\boxed{\text{mem}(x, \text{cons}(x, \ell))} \text{ with } \sigma_2
\end{array}$$

Optionally, a predicate can be declared as a “blocking” predicate, meaning that there are no clauses containing the predicate in the conclusion of the clause. In this case, ProVerif proves properties that hold for *any* definition of the considered blocking predicate.

Semantics of processes. The semantics of processes is defined by the means of a reduction relation $\xrightarrow{\ell}$ between *configurations* which express the current state of the the execution of the processes interacting with the attacker. Formally, a configuration is a triple $\mathcal{E}, \mathcal{P}, \mathcal{A}$ where \mathcal{E} is the set of names used in the configuration, \mathcal{P} is a multiset of processes, and \mathcal{A} is a set of terms representing the knowledge of the attacker.

The full set of rules defining the relation $\xrightarrow{\ell}$ is provided in Appendix (and in [33]) and we only show a small extract below. For example, the following rule represents that an event is triggered:

$$\mathcal{E}, \mathcal{P} \cup \{\{\text{event}(ev).P\}\}, \mathcal{A} \xrightarrow{\text{event}(ev)} \mathcal{E}, \mathcal{P} \cup \{P\}, \mathcal{A}$$

The rule for predicate evaluation (when the predicate evaluates to true) is:

$$\mathcal{E}, \mathcal{P} \cup \{\{\text{let } x_1, \dots, x_n \text{ suchthat } \text{pred} \text{ in } P \text{ else } Q\}\} \rightarrow \mathcal{E}, \mathcal{P} \cup \{P\}, \mathcal{A}$$

when $\text{pred} = p(M_1, \dots, M_k)$ and there exists a substitution σ such that $\text{dom}(\sigma) = \{x_1, \dots, x_n\}$ and $p(M_1\sigma, \dots, M_k\sigma) \in \text{sem}(p)$.

An *execution trace* of a process P is then defined as a sequence of applications of the relation $\xrightarrow{\ell}$ starting from the initial configuration $\mathcal{C}_1 = (\emptyset, \{\{P\}\}, \emptyset)$, i.e. $T = \mathcal{C}_1 \xrightarrow{\ell_1} \dots \xrightarrow{\ell_n} \mathcal{C}_{n+1}$.

In addition to user-defined predicates, ProVerif considers several native predicates: $\text{att}(M)$ indicating that the attacker knows M ; $\text{mess}(M, N)$ indicating that a message N has been sent on the channel M and $\text{event}(ev)$ indicating that an event ev has been raised. Satisfiability of a fact F by a trace T , denoted $T \vdash F$, is given by the labels and configurations in T (e.g. when $F = \text{event}(ev) = \ell_i$ for some i). Naturally, $T \vdash F$ with $\text{pred}(F) \in \mathcal{F}_p$ when $F \in \text{sem}(p)$.

Finally, a correspondence query $F_1 \wedge \dots \wedge F_n \Rightarrow \psi$ can be seen as the first order logic formula $\Psi = \forall \tilde{x}. (F_1 \wedge \dots \wedge F_n \Rightarrow \exists \tilde{y}. \psi)$ where $\tilde{x} = \text{vars}(F_1, \dots, F_n)$ and $\tilde{y} = \text{vars}(\psi) \setminus \tilde{x}$. The correspondence query holds when for all traces T of P , $T \vdash \Psi$.

Lemmas and axioms. The problem that ProVerif tries to solve is undecidable in general [34]; therefore, by design ProVerif is not complete: it may fail to terminate, and it may yield false attacks. Much work has been done to make it more complete in practice. A significant step in this direction introduces *lemmas and axioms* [9] as a

way to guide derivations in ProVerif, and also to deal with some of the abstractions introduced by the tool.

An *axiom* is an instruction to ProVerif to consider some facts as true, even if they cannot be proved by ProVerif from the protocol process. Typically, an axiom is used if one has a separate (perhaps manual) proof of the fact in question. Consider, for example, a smartcard which stores two secrets, s_1 and s_2 . It allows the user to choose either one of them to be revealed, but not both. (This might be used in a lottery, for example.) We could model the smartcard with the following process:

$\text{in}(c, x); (\text{if } x = 1 \text{ then out}(c, s_1) \mid \text{if } x = 2 \text{ then out}(c, s_2))$

The user chooses to enter 1 or 2, and obtains the corresponding secret; after that, the smartcard does not accept any further input.

The intended security property is that at most one secret is revealed. Unfortunately, ProVerif is not able to prove the security of this device with the given process description. The reason is that ProVerif introduces an abstraction, which allows it to consider a derivation in which x has sometimes the value 1 and sometimes the value 2, and hence the output of s_1 and s_2 can both occur. This is not a valid trace; it is a false attack introduced by ProVerif’s abstraction.

Axioms allow us to rectify this situation. We write the process as follows:

$\text{new } st; \text{in}(c, x); \text{event}(\text{Uniq}(st, x));$
 $(\text{if } x = 1 \text{ then out}(c, s_1) \mid \text{if } x = 2 \text{ then out}(c, s_2))$

and we add the axiom

$$\text{Uniq}(st, x_1) \wedge \text{Uniq}(st, x_2) \Rightarrow x_1 = x_2.$$

The axiom asserts that only one value of x is allowed. This axiom is valid, since for a given st there can be only a single input of x . ProVerif is able to use this axiom to prove the security of the device. This kind of axiom, stating that only one value of an input is allowed, is very useful in increasing the precision of ProVerif.

A *lemma* is similar to an axiom as ProVerif uses it in proofs to establish the desired property but it must be able to prove it first (while it does not try to prove axioms). Lemmas are therefore a useful way of decomposing a verification into smaller pieces.

3. A methodology to model protocols with complex data structures

An intuitive, first attempt to model the transparent decryption protocol in ProVerif requires one to define:

- The required equational theories (e.g. public key encryption).
- The predicates and clauses defining the data structure for the log maintainer.
- The predicates and clauses that represent true statements about the proof of presence and extension that the data structure must satisfy.
- The process defining the protocol.
- The security properties that we are interested to test.

Unfortunately, this “monolithic” attempt to prove the security properties does not always work when the protocol has to deal with complex or recursive data structures,

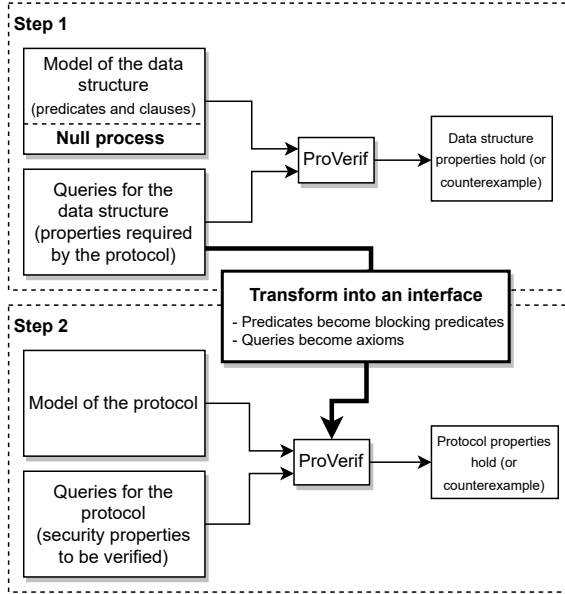


Figure 3. A methodology to prove security in protocols involving complex data structures

and it will often result in ProVerif failing to terminate. In fact, we have encountered this issue even when we try this approach by substituting Merkle trees with a simpler data structure like a hash list.

One way to avoid the monolithic proof is to make an appropriate abstraction of the Merkle tree data structure. This could be achieved by defining the properties the data structure is expected to satisfy, and proving the protocol based on assuming those properties. Later, as a separate step, one can look at whether a particular implementation satisfies the assumed properties.

This observation suggests the following generic methodology to approach proving security when a protocol requires usage of complex data structures. The key idea is to decompose the proof, separately proving security for the data structure and for the protocol separately, using these steps:

- S1: Identify and extract the properties of the data structure that are required by the protocol, and prove that there is an implementation of a data structure that satisfies these properties. As the semantics of the data structure are protocol-independent, so are these properties.
- S2: Assuming that we have a data structure with such properties, prove security for the protocol itself.

We depict this methodology in Fig. 3. For the case of accountable decryption it works as follows. In Step S1, we propose a suitable model for the data structures employed by the ledger, which we instantiate either as a hash list or as a Merkle tree. We also define the desired properties of the data structure from the point of view of the protocol. The reason to consider two data structures is to show that the identified “interface of security properties” and the methodology are sufficiently generic to allocate several data structures and be considered for several protocols. Next, we use ProVerif to prove that the data structure satisfies the properties in the null process.

On the other hand, in Step S2, we adopt the properties

for the data structure as an interface. This requires to take two actions: transform these properties as ProVerif axioms, and transform any related predicate to a blocking predicate. This will allow us to prove the security property of the protocol disregarding the particular implementation of the data structure. Observe that this step is independent from Step S1 above, i.e. they both can be executed in parallel.

By decomposing the problem into two parts, ProVerif has a better chance of avoiding nontermination, without losing soundness of the proof. Moreover the “interface” approach allows for proving security for more complex data structures than Merkle trees, and allows proving security of the protocol for any underlying data structure satisfying the data structure properties.

The remainder of this section describes in detail the methodology, by detailing the implementation of the two data structures commented above, the interface of security properties of the data structure, and the modelling and analysis of the accountable decryption protocol.

3.1. Modeling the ledger data structure

Recall that the ledger is an untrusted party that maintains an append-only log L . From the point of view of proving security properties for a protocol, it does not matter which is the particular data structure employed by the log maintainer, as long as it is append-only, and provides an interface to construct and verify proofs of presence and proofs of extension. Whichever data structure we use must define the clauses from which valid proofs of presence and the proofs of extension can be derived. We discuss two common data structures below, namely hash lists and Merkle trees.

Even though the most interesting case is indeed the latter one, the main ideas of our approach are more clearly seen using hash lists. For this reason, we provide a more detailed explanation on how to model the proofs on hash lists, whereas we provide a brief example on Merkle trees, avoiding routine technicalities that might hinder the main points we want to state. The complete models, both for hash lists and Merkle trees are available at [10].

For convenience, we parameterize the hash function H with two values, e.g. $H(x, h)$, where h is an output of the hash function. This can be implemented simply as concatenating the values x and h in a regular hash function. Hashes of single values x are interpreted as $H(x, h_0)$, where h_0 denotes the null hash.

3.1.1. Hash lists. First, consider the case where L is a hash list. That is, L is represented by h , where

$$h = H(R_n, \dots, H(R_2, H(R_1, h_0)) \dots).$$

In this case, a proof of presence of R_i in the list represented by h simply consists of the elements inserted after R_i , plus the hash of the list before inserting R_i , i.e.

$$\pi = \text{pp}((R_n, \dots, R_{i+1}), h'),$$

where $h' = H(R_{i-1}, \dots, H(R_1, h_0) \dots)$ and pp is the constructor for the proof of presence. The predicate

$\text{verify_pp}(\pi, R_i, h)$ states that π is a valid proof of presence of R_i in the list represented by h . For a hash list, verify_pp will check that the following equation holds:

$$h = H(R_n, \dots, H(R_{i-1}, H(R_i, h')) \dots). \quad (1)$$

In order for ProVerif to handle predicates in the resolution algorithm, we need to provide the set of Horn clauses from which all valid predicates can be derived inductively. In the case of the proof of presence, these clauses are

$$\text{verify_pp}(\text{pp}(\text{nil}, h'), R, H(R, h')), \quad (2)$$

$$\text{verify_pp}(\text{pp}(\ell, h'), R, h) \rightarrow$$

$$\text{verify_pp}(\text{pp}(\text{cons}(Q, \ell), h'), R, H(Q, h)). \quad (3)$$

Clause (2) states the ‘‘base case’’: a proof of presence to verify the last entry on a hash list is an empty list of elements and the immediate hash h' before inserting R . It can be readily seen that the verification from Eq. (1) holds. Clause (3) states the recursive nature of the proof of presence: if a list contains R , an extension of this list with an element Q also contains R , and a proof of presence can be easily derived by prepending Q into the list of elements of the original proof of presence.

For the proof of extension, let L_1 and L_2 be two hash lists represented by h_1 and h_2 , respectively, and with lengths $n_1 \leq n_2$. A proof of extension ρ that L_2 extends L_1 simply consists of the list of elements inserted into L_2 after the last element R_{n_1} inserted into L_1 , that is,

$$\rho = \text{pe}(R_{n_2}, \dots, R_{n_1+1}),$$

where pe is the constructor for proofs of extension. The implementation of predicate $\text{verify_pe}(\rho, h_1, h_2)$, which verifies that ρ is a valid proof of extension for the lists L_1, L_2 , will check that the equation below holds:

$$h_2 = H(R_{n_2}, \dots, H(R_{n_1+1}, h_1) \dots). \quad (4)$$

The predicate verify_pe is defined by the following Horn clauses:

$$\text{verify_pe}(\text{pe}(\text{nil}), h, h), \quad (5)$$

$$\text{verify_pe}(\text{pe}(\ell), h_1, h_2) \rightarrow$$

$$\text{verify_pe}(\text{pe}(\text{cons}(R, \ell)), h_1, H(R, h_2)). \quad (6)$$

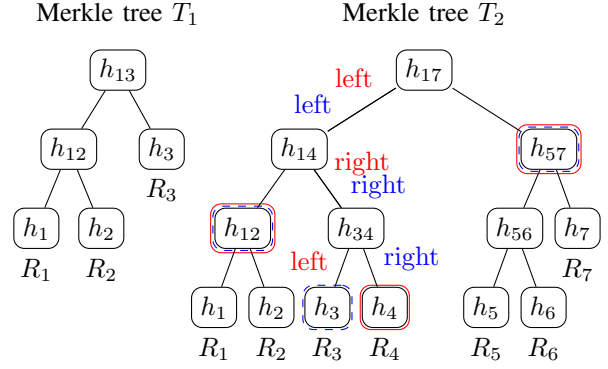
Indeed, Clause (5) above is the ‘‘base case,’’ and indicates that a hash list L represented by h extends itself trivially, hence the proof of extension is the empty list. Clause (6) states that given a valid proof of extension indicating that a list L_2 represented by h_2 extends L_1 represented by h_1 , then $\text{pe}(\text{cons}(R, \ell))$ is a valid proof of extension stating that R appended to L_2 also extends L_1 .

Note that both π and ρ are data structures of size $O(n)$, and the verification of these proofs (Eq. (1) and (4)) also takes time $O(n)$.

Finally, we also require a predicate $\text{repr}(\ell, h)$ to state the fact that h represents a hash list data structure containing the elements in $\ell = (R_n, \dots, R_1)$, inserted in reverse order. Clearly, the clauses defining this predicate are

$$\text{repr}(\text{nil}, h_0),$$

$$\text{repr}(\ell, h) \rightarrow \text{repr}(\text{cons}(R, \ell), H(R, h)).$$



Label h_{ij} is the digest containing R_i, \dots, R_j .

The proof of presence of R_4 in T_2 is $((\text{left}, h_{57}), (\text{right}, h_{12}), (\text{right}, h_{34}))$, shown in blue.

The proof of extension of T_1 to T_2 is $R_3, ((\text{left}, h_{57}), (\text{right}, h_{12}), (\text{left}, h_4))$, shown in red.

Figure 4. Examples of Merkle trees

3.1.2. Merkle trees. We assume that the reader has some familiarity with Merkle trees. It is a more efficient data structure to construct the proofs: the memory and time requirements are reduced to $O(\log n)$ (compared with $O(n)$ for hash lists). We omit most of the formalities and consider the example of Merkle trees depicted in Fig. 4. Recall that a Merkle tree is a binary tree that assigns a hash value for each node, computed as the hash of its child nodes, e.g. $h_{12} = H(h_1, h_2)$. By convention, the hashes on the leaf nodes are defined as the hash of the associated element, i.e. $h_i = H(R_i, h_0)$. Hence, the root tree hash that represents the tree T_2 in Fig. 4 is $h = h_{17} = H(h_{14}, h_{57})$.

A proof of presence that R_i is present in the tree is constructed by providing the complementary node at each tree level. For example, the proof of presence π of R_4 in T_2 is the list of values

$$\pi = \text{pp}((\text{left}, h_{57}), (\text{right}, h_{12}), (\text{right}, h_4)).$$

The constants ‘‘left’’ and ‘‘right’’ are appended to each element to record the relative position of the path leading to the root three hash. Then, $\text{verify_pp}(\pi, R_4, h_{17})$ will check that

$$h = H(H(h_{12}, H(h_3, H(R_4, h_0))), h_{57}).$$

The predicate verify_pp is defined by the following Horn clauses:

$$\text{verify_pp}(\text{pp}(\text{nil}), R, H(R, h_0)), \quad (7)$$

$$\text{verify_pp}(\text{pp}(\ell), R, h_\ell) \rightarrow$$

$$\text{verify_pp}(\text{pp}(\text{cons}((\text{left}, h_r), \ell)), R, H(h_\ell, h_r)), \quad (8)$$

$$\text{verify_pp}(\text{pp}(\ell), R, h_r) \rightarrow$$

$$\text{verify_pp}(\text{pp}(\text{cons}((\text{right}, h_\ell), \ell)), R, H(h_\ell, h_r)). \quad (9)$$

It is not difficult to see that these clauses represent inductively the proof of presence, starting from the base case in Clause (7), and defining recursively the proofs whose next step are left or right paths in Clauses (8) and (9), respectively.

A proof of extension for Merkle trees can be seen as proofs that the last element of the smaller tree is present in the smaller and the larger tree, and a relationship between the two proofs. For example, to prove that T_2 extends T_1 in Fig. 4, the proof of extension ρ must include the following values:

$$\rho = \text{pe}(R_3, ((\text{left}, h_{57}), (\text{right}, h_{12}), (\text{left}, h_4))).$$

The verification of the proof of extension consists of (i) verifying that the list $((\text{left}, h_{57}), (\text{right}, h_{12}), (\text{left}, h_4))$ is a proof of presence of R_3 in T_2 ; and (ii) that this list filtered by only keeping the elements (right, x) is a proof of presence of R_3 in T_1 .

We refer the reader to the repository of the models [10] for the Horn clause definitions of the predicates `repr` and `verify_pe`.

3.2. Properties for the data structure

Regardless of what data structure is used to model the ledger, the protocol expects that a number of properties hold, which can be abstracted away from the particular data structure. In order to express these properties, we consider the same three predicates we presented above (`verify_pp`, `verify_pe` and `repr`) and we axiomatize their intuitive semantics, i.e. `verify_pp` validates proofs of presence, `verify_pe` validates proofs of extension and `repr` validates that a digest represents the contents of the data structure. We present the properties as first-order logical formulas as follows:

Proof for the empty list. A data structure h_0 represents the situation where the list is empty:

$$\forall h. \text{repr}(\text{nil}, h) \implies h = h_0. \quad (\text{P1})$$

Correctness of the proof of presence. `mem`(R, ℓ) iff its presence can be proved:

$$\forall R, \ell, h, \pi.$$

$$\text{repr}(\ell, h) \wedge \text{mem}(R, \ell) \implies \text{verify_pp}(\pi, R, h), \quad (\text{P2})$$

$$\text{repr}(\ell, h) \wedge \text{verify_pp}(\pi, R, h) \implies \text{mem}(R, \ell). \quad (\text{P3})$$

Correctness of the proof of extension. A data structure for a list ℓ extends the data structure for any of its suffixes:

$$\forall R, \ell, h, \rho. \text{repr}(\text{cons}(R, \ell), h) \implies \exists h'. \text{repr}(\ell, h') \wedge \text{verify_pe}(\rho, h', h). \quad (\text{P4})$$

Transitivity of the proof of extension. If h_3 extends h_2 , which in turn extends h_1 , then h_3 extends h_1 :

$$\begin{aligned} &\forall \rho_1, \rho_2, h_1, h_2, h_3. \\ &\text{verify_pe}(\rho_1, h_1, h_2) \wedge \text{verify_pe}(\rho_2, h_2, h_3) \implies \\ &\exists \rho_3. \text{verify_pe}(\rho_3, h_1, h_3). \end{aligned} \quad (\text{P5})$$

Compatibility of the proofs of extension and presence. If an element is present in a list, then it remains present after further elements have been added to it:

$$\begin{aligned} &\forall R, \pi_1, \rho, h_1, h_2. \\ &\text{verify_pp}(\pi_1, R, h_1) \wedge \text{verify_pe}(\rho, h_1, h_2) \implies \\ &\exists \pi_2. \text{verify_pp}(\pi_2, R, h_2). \end{aligned} \quad (\text{P6})$$

```

process
!
new dk; // Decryption key
new sk; // Signing key
out(c, (enc_key(dk), ver_key(sk))); // Output public
keys
new cell; // Memory cell of the trustee
new monitor; // Priv. chan. to store generated
ciphertexts
( out(cell, (0, h0)) // Initialize trustee memory cell
| ( ! // Trustee decrypt protocol
in(cell, (i, h));
in(c, (R, pi, rho, h'));
if verify_pe(rho, h, h') then
if verify_pp(pi, R, h') then
event Decrypted(cell, i, R);
out(c, decr(dk, R));
out(cell, (i+1, h'))
) | ( ! // Trustee sign hash protocol
in(cell, (i, h));
in(c, v);
event Signature(cell, i, h, sign(sk, (v, h)));
out(c, sign(sk, (v, h)));
out(cell, (i, h))
) | ( ! // Subject ciphertext generation
new s;
event Secret(s);
out(c, encr(ek, s));
out(monitor, (s, encr(ek, s)))
) | ( ! // Monitor ciphertext
in(monitor, (s, R));
in(c, =s);
new v;
event Name(v, s, enc_key(dk), ver_key(sk));
out(c, v);
in(c, sigma);
let (=v, h) = checksign(vk, sigma) in
event AfterSeeingSecret(R, h)
) )
) )

```

Figure 5. Transparent decryption model (sketch)

Consistency of digest representation. If two lists are represented by the same digest, then they are equal:

$$\forall \ell_1, \ell_2, h. \text{repr}(\ell_1, h) \wedge \text{repr}(\ell_2, h) \implies \ell_1 = \ell_2. \quad (\text{P7})$$

As described above, (P1)-(P7) can be regarded as an “interface” of security properties, which is instantiated as ProVerif queries when we prove them for the data structure in Step S1 of the methodology, and it is instantiated as ProVerif axioms with blocking predicates when they are used to prove security in the protocol, in Step S2.

3.3. Modelling the transparent-decryption protocol and its security properties

In this section, we provide the most relevant details of the model of the protocol for transparent decryption. A sketch of the model is depicted in Fig. 5. We omit the definitions of the required equational theories for hashing, public key encryption, signatures, types and events. The model essentially consists of four sub-processes running in parallel.

The main process starts by creating private and public keys and initializing the public channel and two private

channels. The requirement of private channels is because, whereas private constant values such as private keys can simply be defined within the scope of the trusted process as new values, the standard way to store mutable values in applied- π calculus is through private channels. Thus, the channels used by our model are:

- Channel `c`: public channel used for communication across the parties.
- Channel `cell`: private channel used as a memory cell for the trustee, in order to store the last seen hash value.
- Channel `monitor`: private channel used to pass secret values and ciphertexts from the subject party to a “monitor process” that will be discussed below.

Below we relate how the four sub-processes map to the different parties presented in Sec. 2.2. However, we remark that, as our trust assumptions are very weak, we do not require to model all the parties so that ProVerif can reason about the security of the protocol:

- Subjects are represented by a sub-process that creates new secrets s and output ciphertexts generated using the trustee’s public encryption key ek . An event `Secret(s)` is required to be used in the definition of the security properties.
- Trustees run the two protocols from Fig. 1, and thus are modelled using two sub-processes. We do not consider threshold decryption schemes for this proof of concept. Note that the first and last step in each of these protocols consists in reading and writing to the trustee memory cell, respectively. Again, the events `Decrypted($cell, i, R$)` and `Signature($cell, i, h, \sigma$)` are used to define the security properties. The former is executed when the trustee accepts to decrypt R , and the latter is used when the trustee casts a signature of its last seen hash.
- The decryption requester and the log L are untrusted entities, as discussed in Sec. 2.2. Hence, there is no need to model them, as the Dolev-Yao adversary will emulate their behaviour.
- Finally, a monitor process is required to ensure that the protocol works as intended and satisfies the claimed security properties. Its main task consists in interacting with the trustee through its second protocol (right column on Fig. 1) by generating a random nonce v and obtaining the signature of the trustee’s last seen hash. Once this happens, the event `AfterSeeingSecret(R, h)` is declared, stating that the monitor has a proof that the trustee has stored the hash value h after the subject has created the ciphertext.

The sub-processes are accordingly replicated to model arbitrary executions of the protocols, and arbitrary number of independent trustees that might monitor different logs.

The formalization of the main security property of the accountable decryption protocol, presented in Sec. 2.2, is

$$\forall R, h. \text{AfterSeeingSecret}(R, h) \implies \exists \pi. \text{verify_pp}(\pi, R, h). \quad (10)$$

Indeed, the location of `AfterSeeingSecret(R, h)` in the monitor process captures the fact that a certain secret s associated to ciphertext (decryption request) R has been observed in the public channel, and the last

observed hash value by the trustee is h . Therefore, any occurrence of this event means that there must exist a proof of presence π stating that R is in the data structure represented by h .

In order to prove this query, we consider some additional lemmas such as the following one:

$$\begin{aligned} &\forall v, s, i, h, cell, sk, dk. \\ &\text{Signature}(cell, i, h, \text{sign}(sk, (v, h))) \wedge \\ &\text{Name}(v, s, \text{enc_key}(dk), \text{ver_key}(sk)) \implies \\ &\exists j. \text{Decrypted}(cell, j, \text{encr}(\text{enc_key}(dk), s)) \wedge j < i. \end{aligned}$$

Notice that the memory cell of the trustee stores two pieces of information: the number of times it decrypted a ciphertext, and the latest hash value it received. Hence this lemma states that when the trustee signed the name v generated by the monitor (which monitors the trustee with encryption key $\text{enc_key}(dk)$ and verification key $\text{ver_key}(sk)$) after receiving a secret s , then the trustee must have decrypted it *strictly before* (i.e. $j < i$).

This lemma allows us to help ProVerif by linking the content of the memory cell (i.e. the number of times it decrypted an encryption) with the order of events that were emitted in the trace. Such links are usually abstracted away by ProVerif during the saturation procedure hence the need for us to provide it within a lemma. Note that we also considered some additional lemmas and axioms, specific to the management of memory cells, in the vein of [35].

Our ProVerif models only take couple of seconds to execute on a standard laptop.

4. Extending ProVerif to support arbitrary predicates in lemmas and axioms

To complete the methodology depicted in Fig. 3 on our running example, we need to (1) prove the properties **P1** to **P7**, and (2) prove the main protocol while expressing the properties **P1** to **P7** as axioms. All the properties in our interface are in fact correspondence properties that are within the scope of ProVerif. For example, the query corresponding to Property **P6** would be expressed as follows:

```
query pe1, pe2, pe3: proof_of_extension,
      d1, d2, d3: digest;
verify_pe(pe1, d1, d2) && verify_pe(pe2, d2, d3)
  => verify_pe(pe3, d1, d3)
```

However, in its current version, ProVerif imposes a syntactic restriction on axioms and lemmas: the conclusion of a lemma can only contain events, equalities, disequalities and blocking user-defined predicates. Specifically, the native facts $\text{att}(M)$ and $\text{mess}(M, N)$ as well as the clause-based user-defined predicates cannot be used in the conclusion of a lemma.

This restriction of facts prevents us from achieving both steps of our methodology. The second step is unattainable as ProVerif will directly reject such a query if it is written as an axiom. For the first step, the query will be accepted by ProVerif but it will fail to prove it. Such a query requires a proof by induction, which internally corresponds to transforming a query into an *inductive lemma* that has the same syntactic limitation as declared axioms and lemmas. By extending ProVerif to allow any

predicates in the conclusion of axioms and lemmas, we are able to complete our methodology.

In this section, we provide a high-level description of ProVerif's procedure and how we extended it.

4.1. Description of ProVerif's procedure

Horn clauses are used to describe the semantics of user-defined predicates as previously described but they are also the building blocks of ProVerif's internal procedure to prove a secrecy property and more generally a correspondence property. Specifically, ProVerif first translates the protocol given as input into a set \mathbb{C} of Horn clauses. It then proceeds to *saturate* this set \mathbb{C} and the clauses that define the user-defined predicates, yielding a simpler set of clauses that derives the same facts. The procedure completes by verifying that the saturated clauses satisfy the security property.

Translation into Horn clauses. In addition to the user-defined predicates, ProVerif considers natively four additional predicates over terms representing the interactions between the attacks and the processes: $\text{att}(M)$, $\text{mess}(M, N)$, and two predicates for events $\text{s-event}(ev)$ and $\text{event}(ev)$. *Sure-events*, i.e. $\text{s-event}(ev)$, will only appear in hypotheses of Horn clauses, whereas events $\text{event}(ev)$ will only appear in their conclusion. This separation ensures that events are not *resolved* during the saturation procedure and so their occurrence in a Horn clause is preserved through resolution rule.

Using these predicates, ProVerif generates a set of clauses representing the capabilities of the attacker, which include, for example:

$$\text{att}(x) \wedge \text{att}(y) \rightarrow \text{att}(\text{encr}(x, y)) \quad (11)$$

$$\text{att}(\text{encr}(x, y)) \wedge \text{att}(x) \rightarrow \text{att}(y) \quad (12)$$

$$\text{att}(x) \wedge \text{mess}(x, y) \rightarrow \text{att}(y) \quad (13)$$

$$\text{att}(x) \wedge \text{att}(y) \rightarrow \text{mess}(x, y) \quad (14)$$

The first two clauses model that the attacker can encrypt and decrypt provided that it knows the secret key. The last two clauses model that the attacker can read and write on a channel that it knows.

The formal description of processes into Horn clauses is out of scope of this paper (see [9] for more details), but we provide some intuition in the following example.

Example 2. *The translation of the Ciphertext Generator process would yield at least the clause*

$$\text{s-event}(\text{Secret}(s)) \rightarrow \text{att}(\text{encr}(ek, s)) \quad (15)$$

indicating that the attacker can obtain the encryption of the secret s by the key ek . The Horn clause also indicates that the event $\text{Secret}(s)$ is triggered before the encryption is sent to the attacker.

Similarly, the translation of process modelling the accountable decryption device will generate, in particular, the clause:

$$\begin{aligned} &\text{mess}(d, (i, H_0)) \wedge \text{att}(\text{encr}(ek, x)) \wedge \text{att}((pi, r, H_1)) \wedge \\ &\text{verify_pe}(r, H_0, H_1) \wedge \text{verify_pp}(pi, \text{encr}(ek, x), H_1) \wedge \\ &\text{s-event}(\text{Decrypted}(i, \text{encr}(ek, x))) \rightarrow \text{att}(x) \quad (16) \end{aligned}$$




Figure 6. A schematic summary of ProVerif's saturation procedure

This clause represents the informal statement: Assuming that the cell of the device contains the i -th digest received H_0 , if the attacker can provide a proof of extension r from H_0 to a new digest H_1 , a proof of presence pi of some ciphertext $\text{encr}(ek, x)$ in the digest H_1 then the attacker can obtain the plain text x . As in the previous clause, the event $\text{Decrypted}(i, \text{encr}(ek, x))$ will be triggered before the attacker obtains x .

Saturation. The core step in the saturation procedure consists of taking two existing Horn clauses and combining them into a new one, hopefully simpler. This process is called the *resolution step*. For example, the hypothesis of clause 16 contains in hypothesis the fact $F = \text{att}(\text{encr}(ek, x))$ representing that the attacker must know the ciphertext $\text{encr}(ek, x)$. In order to deduce all the possible ways the attacker may deduce this ciphertext, ProVerif resolves this fact by combining it with Horn clauses whose conclusion can be unified with F . This is the case with clause 15, which results in the following clause:

$$\begin{aligned} &\text{mess}(d, (i, H_0)) \wedge \text{s-event}(\text{Secret}(s)) \wedge \text{att}((pi, r, H_1)) \wedge \\ &\text{verify_pe}(r, H_0, H_1) \wedge \text{verify_pp}(pi, \text{encr}(ek, s), H_1) \wedge \\ &\text{s-event}(\text{Decrypted}(i, \text{encr}(ek, s))) \rightarrow \text{att}(s). \end{aligned}$$

Note that the hypothesis $\text{s-event}(\text{Secret}(s))$ of clause 15 has replaced F in the new clause and the unification of the two facts resulted in x being instantiated by s .

In our example, the fact F in clause 16 could also have been resolved with clause 12, which would result in a new rule like clause 16 but with F replaced by the two facts $\text{att}(\text{encr}(x', \text{encr}(ek, x)))$ and $\text{att}(x')$. This would directly lead to a loop during the saturation procedure as the former fact could once again be resolved by the same clause 12. To avoid this problem, ProVerif uses a *selection function* on Horn clauses that returns the set of facts from the hypothesis of a clause that can be resolved, i.e. $\text{sel}(H \rightarrow C) = S$ with $S \subseteq H$. In particular, the facts $\text{att}(x)$ with x a variable are never in S .

The resolution process repeats until a fixed point is reached, i.e. until no resolution can produce a clause that is not redundant with an existing one. Once a fixed point is reached, ProVerif only keeps the set of *saturated clauses*, which are the clauses R from \mathbb{C} such that $\text{sel}(R) = \emptyset$. A schematic summary of ProVerif's saturation procedure is given in Fig. 6.

Generalizing the events and sure-events separation. We previously mentioned that ProVerif considers two

different predicates for events: the sure-event predicate s-event and the event predicate event. During the translation of processes into Horn clauses, the former type of event only occurs in the hypotheses of the clauses whereas the latter only occurs in the conclusion of the clauses. It ensures that an application of the resolution rule never resolves a sure-event s-event(ev) from a clause. This is critical for the verification of correspondence queries. For example, to verify the query $event(A) \Rightarrow event(B)$, ProVerif will check that for all saturated clauses in \mathbb{C}_{sat} of the form $H \rightarrow event(A)$, the sure-event s-event(B) occurs in H . In this case, ProVerif concludes that the query holds. Note that to ensure soundness, the selection function will never consider sure-events, i.e. $s-event(ev) \notin sel(R)$, which guarantees that the resolution rule does not try to resolve a fact that cannot be resolved by design.

Since sure-events are never selected by the selection function, they can be seen as the *blocking* counterpart of the predicate event. To add clause-based, user-defined predicates in the conclusion of lemmas, as well as the predicates $att(M)$ and $mess(M, N)$, we generalize this concept by associating to all predicates p a blocking predicate denoted $b-p$. Therefore, we consider natively the predicates $b-att$ and $b-mess$. Moreover, for all user-defined predicates $p \in \mathcal{F}_p$, we consider the blocking predicate $b-p$ and denote by \mathcal{F}_{bp} their set. Finally, we rename the predicate s-event as b-event.

We also amend the selection function by additionally requiring that no blocking fact can be selected: For all F , $b-F \notin sel(H \rightarrow C)$. Thus, the resolution rule will never attempt to resolve a blocking predicate.

Application of lemmas. ProVerif axioms and proved lemmas are applied during the simplification phase of the saturation procedure. Consider a lemma $\bigwedge_{i=1}^n F_i \Rightarrow \bigvee_{j=1}^m \psi_j$, where ψ_j are conjunctions of facts. In the current ProVerif, each ψ_j could only be composed of events, disequalities and inequalities. The lemma would be applied on a clause $H \rightarrow C$ when there exists a substitution σ such that $F_i\sigma$ is in H for all $i \in \{1, \dots, n\}$. The application of the lemma would then produce a set of m clauses $\{H \wedge \psi'_j\sigma \rightarrow C\}_{j=1}^m$, where each ψ'_j is the disjunct ψ_j with events replaced by their sure-event counterpart.

Thanks to our extension, the disjunct ψ_j may now contain any type of predicate. We update the simplification rule by first defining the transformation $\lceil \psi \rceil^b$ built from ψ where all facts $p(M_1, \dots, M_n)$ in ψ are replaced by $b-p(M_1, \dots, M_n)$. The rule for applying lemmas is then defined as follows:

$$\frac{\mathbb{C} \cup \{R = (H \rightarrow C)\} \quad (\bigwedge_{i=1}^n F_i \Rightarrow \bigvee_{j=1}^m \psi_j) \in \mathcal{L} \quad \forall i, b-F_i\sigma \in H \text{ or } F_i\sigma \in H}{\mathbb{C} \cup \{H \wedge \lceil \psi_j \sigma \rceil^b \rightarrow C\}_{j=1}^m}$$

Note that the application condition requires that either $b-F_i\sigma \in H$ or $F_i\sigma \in H$. Indeed, the lemma can be applied on facts introduced by a previous application of another lemma or more commonly on an event, hence the condition $b-F_i\sigma \in H$. Since a lemma may have an attacker fact, a message fact or a fact using a user predicate in its premisses, we also need to match the non-blocking

form of the fact, i.e. $F_i\sigma \in H$. Note that all facts added in the clauses, i.e. facts in $\lceil \psi_j \sigma \rceil^b$, are blocking.

Improving other simplification rules. The purpose of applying lemmas is to increase ProVerif precision or to help it terminate. As such, we can amend other simplification rules used by ProVerif to benefit from the blocking facts. For example, ProVerif employs the following simplification rule to remove tautologies:

$$\frac{\mathbb{C} \cup \{F \wedge H \rightarrow F\}}{\mathbb{C}}$$

Now that we may have blocking predicates in the hypotheses of the clause, we can consider an additional tautology simplification rules defined as follows:

$$\frac{\mathbb{C} \cup \{b-F \wedge H \rightarrow F\} \quad \text{pred}(p) \notin \mathcal{F}_p}{\mathbb{C}}$$

We also improve the simplification rule that removes the redundant facts from the hypotheses of a clause, which is critical to avoid termination issues:

$$\frac{\mathbb{C} \cup \{H' \wedge H \wedge \phi \rightarrow C\} \quad \lceil H'\sigma \rceil^b \subseteq \lceil H \rceil^b \quad \phi \models \phi\sigma \quad \text{dom}(\sigma) \cap \text{vars}(H, C) = \emptyset}{\mathbb{C} \cup \{H \wedge \phi\sigma \rightarrow C\}}$$

Intuitively, this rule states that to derive C , it suffices to know the derivations of H , as a derivation for H' can be build from the derivations of $H'\sigma \subseteq H$. Note that by requiring $\lceil H'\sigma \rceil^b \subseteq \lceil H \rceil^b$, a clause $att(M) \wedge b-att(M) \wedge G \rightarrow C$ can be either simplified into $att(M) \wedge G \rightarrow C$ or into $b-att(M) \wedge G \rightarrow C$. Both variants are correct and their application may be parametrized depending on one's needs. The former could be used to discard the application of a lemma that is redundant with the hypothesis of the clause whereas the latter could help terminate, as the fact $b-att(M)$ will not be resolved.

Verification of the query. Once the saturation process ends, ProVerif still needs to verify the queries. Intuitively, on a query $F_1 \wedge \dots \wedge F_n \Rightarrow \psi$ with $F_i = p_i(t_1^i, \dots, t_{m_i}^i)$, ProVerif starts by generating a clause $F_1 \wedge \dots \wedge F_n \rightarrow C$ with $C = q(t_1^1, \dots, t_{m_n}^n)$ and q a special predicate only used in the verification of query. Typically, C represents the conjunction of facts F_1, \dots, F_n . ProVerif then applies once again a saturation of $\mathbb{C}_{sat} \cup \{F_1 \wedge \dots \wedge F_n \rightarrow C\}$ and checks the validity of ψ on the obtained saturated set. Note that this second saturation is slightly different from the first saturation in the sense that it is *order preserving*, i.e. in a clause $F \wedge H \rightarrow q(t_1^1, \dots, t_{m_n}^n)\sigma$, ProVerif can indicate, for any i , whether the fact F was generated when resolving F_i . This property is particularly important for proving queries by induction. The details of how the orders are preserved can be found in [9].

Example 3. Let us show how we are able to prove by induction the transitivity of proof of extension in the hash list data structure:

```
query pe1, pe2, pe3:proof_of_extension,
      d1, d2, d3:digest;
verify_pe(pe1, d1, d2) && verify_pe(pe2, d2, d3)
  => verify_pe(pe3, d1, d3) [induction].
```


To ease the reading, let us denote by v_{pe} the predicate `verify_pe`. Recall that the proof of extension is defined by the following two clauses.

$$v_{pe}(pe(\text{nil}), d, d), \quad (17)$$

$$v_{pe}(pe(\ell), d_1, d_2) \rightarrow v_{pe}(pe(\text{cons}(x, \ell)), d_1, H(x, d_2)). \quad (18)$$

These clauses will be left unchanged through the saturation procedure, i.e. if $\mathbb{C} = \{(17), (18)\}$ then $\text{sat}(\mathbb{C}) = \mathbb{C}$.

To verify the query, ProVerif first considers the query clause $RQ = F_1 \wedge F_2 \rightarrow C$ where $F_1 = v_{pe}(pe_1, d_1, d_2)$, $F_2 = v_{pe}(pe_2, d_2, d_3)$ and $C = q(pe_1, d_1, d_2, pe_2, d_2, d_3)$. ProVerif then applies the saturation procedure on $\mathbb{C} \cup \{RQ\}$.

We illustrate the application of the inductive lemma on one of the resolutions that ProVerif will apply during this second saturation procedure: the resolution of $v_{pe}(pe_2, d_2, d_3)$ from RQ with the clause (18) which yields the following clause

$$RQ' = F_1\sigma \wedge v_{pe}(pe(\ell), d_2, d'_3) \rightarrow C\sigma$$

with $\sigma = \{pe_2 \mapsto pe(\text{cons}(x, \ell)), d_3 \mapsto H(x, d'_3)\}$.

Let us denote $F_3 = v_{pe}(pe(\ell), d_2, d'_3)$. Note that $C\sigma$ represents the conjunction $F_1\sigma \wedge F_2\sigma$ with:

$$F_2\sigma = v_{pe}(pe(\text{cons}(x, \ell)), d_2, H(x, d'_3))$$

Since F_3 was obtained while resolving $F_2\sigma$, any instantiation $F_3\alpha$ would be satisfied strictly before $F_2\sigma\alpha$ (we detail this notion in the next section). Thus, ProVerif can apply our inductive hypothesis on $F_1\sigma, F_3$ and so it will add the blocking fact $b\text{-}v_{pe}(pe', d_1, d'_3)$ in the hypothesis of the clause, yielding:

$$RQ_2 = F_1\sigma \wedge v_{pe}(pe(\ell), d_2, d'_3) \wedge b\text{-}v_{pe}(pe', d_1, d'_3) \rightarrow C\sigma$$

On the clause RQ_2 , ProVerif will be able prove the query, i.e. finding a derivation of $v_{pe}(pe_3, d_1, d_3)\sigma = v_{pe}(pe_3, d_1, H(x, d'_3))$ for some pe_3 , by using the blocking fact $b\text{-}v_{pe}(pe', d_1, d'_3)$ and the clause (18).

4.2. Soundness

The soundness of the saturation procedure comes in three steps.

Derivation of satisfiable facts. Given a process P and a trace T of P , one can show that for all satisfiable facts F in T , i.e. $T \vdash F$, there exists a derivation \mathcal{D} of F in the set of initial clauses $\mathbb{C}_{init}(P)$ translated from P and the set of clauses $\mathbb{C}_b(T) = \{\lceil F \rceil^b \mid T \vdash F\}$. The original soundness result [9, Theorem 1] was only considering the sure-events, i.e. $\{\text{s-event}(ev) \mid T \vdash \text{event}(ev)\}$.

Since blocking predicates are an artifice to prevent resolution within the saturation procedure, they should not affect their satisfiability in a trace T . Hence, we augment the satisfaction relation \vdash by requiring that for all blocking facts $b\text{-}F$, $T \vdash b\text{-}F$ if and only if $T \vdash F$.

In the current ProVerif, the derivation \mathcal{D} is shown to satisfy an invariant on the Horn clauses labeling its nodes that relates satisfaction of facts with the size of the trace. Consider a ground fact F such that $T \vdash F$. We define $\min_T(F) = \min\{|T'| \mid T' \text{ prefix of } T \wedge T' \vdash F\}$ which represents the size of the smallest prefix of T that satisfies

F . The invariant intuitively indicates that in a derivation, all instantiated facts are satisfied by T and if a clause $H \rightarrow C$ is used to derive C then the facts in H must be satisfied strictly before C in the trace T . We extend this invariant to take into account the user-defined predicates and blocking predicates

However, when a fact F corresponds to a user-defined predicates, e.g. $F = p(M_1, \dots, M_n)$, we always have $\min_T(F) = 0$ as the satisfiability of this predicate do not depend on the protocol but only on the value of the terms M_1, \dots, M_n . Hence, $T \vdash F$ if and only if F is satisfied in the empty trace. This prevents us from ordering the user-defined predicates within a clause. Hence, we strengthen the minimality function on user-defined predicates by considering the size of the derivation of $p(M_1, \dots, M_n)$ in \mathbb{C}_{user} instead of its satisfiability in T . Formally, we consider $\min_{T, \mathbb{C}_{user}}(F)$ defined as:

- $\min_{T, \mathbb{C}_{user}}(F) = \min_T(F)$ when $\text{pred}(F) \notin \mathcal{F}_p \cup \mathcal{F}_{bp}$
- $\min_{T, \mathbb{C}_{user}}(F) = \min\{|\mathcal{D}| \mid \mathcal{D} \text{ derives } F' \text{ from } \mathbb{C}_{user}\}$ when $F = F'$ or $F = b\text{-}F'$ with $\text{pred}(F') \in \mathcal{F}_p$.

We can show state the main invariant on derivations.

Invariant 1. Let T be a trace. Let \mathbb{C}_{user} a set of clauses defining predicates in \mathcal{F}_p . Let \mathcal{D} a derivation. We say that the invariant $\text{Inv}_{T, \mathbb{C}_{user}}(\mathcal{D})$ holds when for all nodes of \mathcal{D} labeled by a Horn clause $R = H \rightarrow C$ and a substitution σ , $T \vdash C$ and for all $F \in H$, $T \vdash F$. Moreover,

- if R is the attacker rule (13) then $\min_T(\text{att}(x)\sigma) < \min_T(\text{att}(y)\sigma)$ and $\min_T(\text{mess}(x, y)\sigma) = \min_T(\text{att}(y)\sigma)$;
- otherwise for all $F \in H$, if $(\text{pred}(C) \notin \mathcal{F}_p$ and $\text{pred}(F) \notin \mathcal{F}_p \cup \mathcal{F}_{bp}$) or $(\text{pred}(C) \in \mathcal{F}_p$ and $\text{pred}(F) \in \mathcal{F}_p)$ then $\min_{T, \mathbb{C}_{user}}(F\sigma) < \min_{T, \mathbb{C}_{user}}(C\sigma)$

As previously mentioned, the invariant intuitively states that all facts in the derivation are satisfied in T . The first bullet point indicates that in the case of the attacker rule (13), representing that the attacker may listen on a channel if it can deduce it, the two facts $\text{mess}(x, y)$ and $\text{att}(y)$ are satisfied at the same moment on T . The second bullet point indicates that for all other rules, the hypotheses are satisfied strictly before the conclusion, in the sense of $\min_{T, \mathbb{C}_{user}}(\cdot)$ and if they belong to the same category, i.e. standard or user-defined predicates.

Preservation of the invariant during the saturation.

In the second step of the soundness proof, one can show that one round *resolution-simplification-filtering* preserves the derivability of facts, i.e. if \mathbb{C}' is the new set of clauses after a round of *resolution-simplification-filtering* on \mathbb{C} , then F derivable in \mathbb{C} implies that F is derivable in \mathbb{C}' , with both derivations satisfying Invariant 1.

The soundness of the new tautology rule is directly given by Invariant 1. Indeed, if $\text{pred}(F) \notin \mathcal{F}_p$ then $\text{pred}(b\text{-}F) \notin \mathcal{F}_p \cup \mathcal{F}_{bp}$ and so $\min_{T, \mathbb{C}_{user}}(b\text{-}F) < \min_{T, \mathbb{C}_{user}}(F)$ which contradicts minimality.

Note that the soundness of the new rule for redundant facts is direct from the fact that invariant 1 is stable by removal of hypotheses.

Finally, we show that the application of lemmas is also sound. This is achieved thanks again to Invariant 1. In particular, when a lemma $\bigwedge_{i=1}^n F_i \Rightarrow \bigvee_{j=1}^m \psi_j$ matches

H , i.e. $F_i\sigma$ is in H for all $i \in \{1, \dots, n\}$, it implies that each $F_i\sigma$ are satisfied in T . Since the lemma holds on T , at least one of the $\psi_j\sigma$ also holds in T and so all facts in $[\psi_j\sigma]^b$ are derivable from $\mathbb{C}_b(T)$.

Application of inductive lemmas. In [9], a query $F_1 \wedge \dots \wedge F_n \Rightarrow \psi$ may be proved by induction on the multiset $\{\{\min_T(F_1\sigma), \dots, \min_T(F_n\sigma)\}\}$ when $T \vdash F_i\sigma$ for all $i = 1, \dots, n$. In practice, ProVerif transforms the query into an *inductive lemma* which is typically a lemma implied by the query. This lemma will be applied as a normal lemma during the saturation procedure except on the attacker rule (I3).

The main soundness argument for applying inductive lemma during the saturation procedure is as follows. As we prove the query by induction on $\mathcal{M} = \{\{\min_T(F_1\sigma), \dots, \min_T(F_n\sigma)\}\}$, we can assume that the query holds for any instance of the premise $F_1\sigma', \dots, F_n\sigma'$ such that $\{\{\min_T(F_1\sigma'), \dots, \min_T(F_n\sigma')\}\} < \mathcal{M}$. Now consider a rule $H \rightarrow C$ used in a derivation of $F_1\sigma, \dots, F_n\sigma$, we know from Invariant 1 that all facts F in H are satisfied strictly before C which is itself satisfied strictly before at least one of the $F_i\sigma$, i.e. $\min_T(F) < \min_T(F_i\sigma)$. Therefore, if $F_1\sigma', \dots, F_n\sigma' \in H$ then we deduce that $\{\{\min_T(F_1\sigma'), \dots, \min_T(F_n\sigma')\}\} < \mathcal{M}$ and so we know that the conclusion of the inductive lemma holds, which allows us to apply it.

With user-defined predicates, we cannot use $\min_T(\cdot)$ because the satisfiability of a user-defined predicate does not depend on the trace but on the set of Horn clauses \mathbb{C}_{user} given as input. Instead, we base the induction on the minimal size of the derivation of F_i in \mathbb{C}_{user} , i.e. $\min_{T, \mathbb{C}_{user}}(F_i\sigma)$. More specifically, if for all $i = 1, \dots, n$, $\text{pred}(F_i)$ is a user-defined predicate, we prove the query by induction on the multiset $\{\{\min_{T, \mathbb{C}_{user}}(F_i\sigma)\}_{i=1}^n\}$.

To prove a query by induction whose premises contain both user-defined predicates and standard predicates, i.e. att and mess , we combine the two inductive measures into a single one: without loss of generality, consider the query $F_1 \wedge \dots \wedge F_n \Rightarrow \psi$ where $\text{pred}(F_1), \dots, \text{pred}(F_{k-1})$ are standard predicates and $\text{pred}(F_k), \dots, \text{pred}(F_n)$ are user-defined predicates. We prove the query for all traces T and all substitutions σ with $T \vdash F_i\sigma$, for all $i = 1, \dots, n$, by induction using the lexicographic order on $(\{\{\min_{T, \mathbb{C}_{user}}(F_i\sigma)\}_{i=1}^{k-1}, \{\{\min_{T, \mathbb{C}_{user}}(F_i\sigma)\}_{i=1}^n\})$.

Restriction to \mathbb{C}_{sat} . The third and final step consists of showing that when the fix point is reached, restricting the \mathbb{C} to \mathbb{C}_{sat} also preserves derivability. This step remains unchanged from [9, Theorem 2].

The general soundness of the saturation procedure is given by the following property.

Theorem 1. *Let P a process and T a trace of P . Let \mathcal{L} a set of lemmas that hold on P . Let \mathcal{L}_i be a set of inductive lemmas. We denote by $\mathbb{C}_{init}(P)$ the set of initial clauses translated from P . Let F be a fact such that $T \vdash F$. Let $\mathcal{M} = (\emptyset, \min_{T, \mathbb{C}_{user}}(F))$ when $\text{pred}(F) \in \mathcal{F}_p$ and $\mathcal{M} = (\min_{T, \mathbb{C}_{user}}(F), \emptyset)$ otherwise.*

If the lemmas in \mathcal{L}_i hold up to \mathcal{M} excluded then there exists a derivation \mathcal{D} of F from $\text{sat}(\mathbb{C}_{init}(P)) \cup \mathbb{C}_b(T)$ such that $\text{Inv}_{T, \mathbb{C}_{user}}(\mathcal{D})$.

This theorem is a simplified version of [9, Theorem 1 and 2, simplified] adapted to user-defined predicates and generalised lemmas. In particular, the key difference between these two soundness results is that we consider in the set $\mathbb{C}_b(T)$ the blocking counter part of all satisfiable facts by the trace T . Note that $\mathbb{C}_b(T)$ also includes the blocking counter part of all true instances of predicates $p(M_1, \dots, M_n)$ where p is a user-defined predicate.

4.3. Applications

As we previously showed, generalizing lemmas, axioms and inductive proofs is paramount for our new methodology. Thanks to the inductive proofs, we are able to show with ProVerif all the properties P1 to P7 for both hash list and Merkle tree data structures. Moreover, as illustrated in Section 3.3, we are also able to prove the security of certificate transparency and transparent decryption with all properties P1 to P7 declared as axioms. All the proofs of the properties in the interface are proved in a single file for hash list in less than a second. For Merkle trees, the proofs are separated in five different files, each taking less than a second to be verified. The proofs of the protocols themselves with the interface are also done in less than two seconds.

Note that the generalisation of lemmas and axioms can also be of use outside of our methodology. For instance, an earlier version of our work has already been used to prove the Encrypted Client Hello extension of the TLS protocol [36]. Specifically, the saturation procedure was entering into a loop when trying to resolve a clause of the form $\text{psk}(id, k, c, s) \wedge H \rightarrow \text{att}(id)$ where $\text{psk}(id, k, c, s)$ intuitively represented some shared key k with identity id between a server s and a client c . To prevent the loop, they added a lemma indicating that when $\text{psk}(id, k, c, s)$ holds then the identity was already known to the attacker or else was honestly generated:

```
lemma id, k, c, s:bitstring;
psk(id, k, c, s) ==> attacker(id) |
id = honest_id(s, c, k) [induction].
```

The application of the inductive lemma on $\text{psk}(id, k, c, s) \wedge H \rightarrow \text{att}(id)$ would yield two clauses, one with $\text{b-att}(id)$ in the hypothesis and one where id is instantiated by $\text{honest}_{id}(s, c, k)$. The former would be removed by our new tautology rule, whereas the instantiated second clause would avoid the loop during the saturation.

5. Conclusion

Transparency in security protocols plays a fundamental role in minimizing the trust conditions for the parties involved. A clear example is that of certificate transparency, where strong security and trust assumptions on certificate authorities can be relaxed by requiring that they publish certificate issuances on a public ledger. In many cases, transparency might well constitute a required building block expected by users of a certain service, especially when it involves incursions into their privacy for a variety of reasons, a situation which is handled by transparent decryption protocols.

For these reasons, it is important to properly verify the core security property of transparency protocols, e.g. that only legitimate certificates are produced (certificate transparency) or that decryptions only take place if the requests for them are entered in a public ledger, i.e. visible to users (transparent decryption).

Because of the complex data structures that transparency protocols rely on, verifying their properties has led to designing a proof-decomposition methodology for ProVerif, as well as adding new features to the way lemmas and axioms are handled. We expect that our methodology and ProVerif enhancements can be applied to other kind of protocols involving tree-based data structures (binary trees, radix trees) and perhaps also other kinds of data structures (e.g. Bloom filters).

Acknowledgements

This work received funding from EPSRC projects *CAP-TEE: Capability Architectures for Trusted Execution*; *SIPP: Secure IoT Processor Platform with Remote Attestation*, and *User-controlled hardware security anchors: evaluation and designs*. It also received funding from the France 2030 program managed by the French National Research Agency under grant agreement No. ANR-22-PECY-0006.

References

- [1] B. Laurie, “Certificate transparency,” *Communications of the ACM*, vol. 57, no. 10, pp. 40–46, 2014.
- [2] B. Laurie, E. Messeri, and R. Stradling, “Certificate transparency version 2.0,” Internet Requests for Comments, RFC Editor, RFC 9162, Dec. 2021. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc1654.txt>
- [3] M. D. Ryan, “Making decryption accountable,” in *Cambridge International Workshop on Security Protocols*. Springer, 2017, pp. 93–98.
- [4] A. Armando, D. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuéllar, P. H. Drielsma, P.-C. Héam, O. Kouchnarenko, J. Mantovani *et al.*, “The avispa tool for the automated validation of internet security protocols and applications,” in *International conference on computer aided verification*. Springer, 2005, pp. 281–285.
- [5] V. Cheval, S. Kremer, and I. Rakotonirina, “Deepsec: deciding equivalence properties in security protocols theory and practice,” in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 529–546.
- [6] R. Chadha, V. Cheval, Ș. Ciobăcă, and S. Kremer, “Automated verification of equivalence properties of cryptographic protocols,” *ACM Transactions on Computational Logic (TOCL)*, vol. 17, no. 4, pp. 1–32, 2016.
- [7] S. Meier, B. Schmidt, C. Cremers, and D. Basin, “The tamarin prover for the symbolic analysis of security protocols,” in *International conference on computer aided verification*. Springer, 2013, pp. 696–701.
- [8] B. Blanchet, “Modeling and verifying security protocols with the applied pi calculus and ProVerif,” *Foundations and Trends in Privacy and Security*, vol. 1, no. 1–2, pp. 1–135, Oct. 2016.
- [9] B. Blanchet, V. Cheval, and V. Cortier, “Proverif with lemmas, induction, fast subsumption, and much more,” in *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022*. IEEE, 2022, pp. 69–86.
- [10] “Proverif models and proverif source code,” <https://www.dropbox.com/sh/gbn5dy0amz1106f/AACbcILzg8o1Bhf5D3nMFa2Wa?dl=0>, 2022.
- [11] D. Basin, C. Cremers, T. H.-J. Kim, A. Perrig, R. Sasse, and P. Szalachowski, “ARPKI: Attack resilient public-key infrastructure,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014, pp. 382–393.
- [12] J. Yu, V. Cheval, and M. Ryan, “DTKI: A new formalized pki with verifiable trusted parties,” *The Computer Journal*, vol. 59, no. 11, pp. 1695–1713, 2016.
- [13] M. D. Ryan, “Enhanced certificate transparency and end-to-end encrypted mail,” in *NDSS Symposium*, 2014.
- [14] R. Künnemann, I. Esiyok, and M. Backes, “Automated verification of accountability in security protocols,” in *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*. IEEE, 2019, pp. 397–39716.
- [15] J. Kroll, E. Felten, and D. Boneh, “Secure protocols for accountable warrant execution,” 2014, https://www.jkroll.com/papers/warrant_paper.pdf.
- [16] D. Nuñez, I. Agudo, and J. Lopez, “Escrowed decryption protocols for lawful interception of encrypted data,” *IET Information Security*, vol. 13, no. 5, pp. 498–507, 2019.
- [17] L. Idan and J. Feigenbaum, “Prshare: A framework for privacy-preserving, interorganizational data sharing,” in *Proceedings of the 19th Workshop on Privacy in the Electronic Society*, 2020, pp. 137–149.
- [18] M. Li, Y. Chen, C. Lal, M. Conti, M. Alazab, and D. Hu, “Eunomia: Anonymous and secure vehicular digital forensics based on blockchain,” *IEEE Transactions on Dependable and Secure Computing*, 2021.
- [19] “Privacy-preserving accountable decryption,” <https://pad.tech>, 2022.
- [20] J. A. Kroll, “Accountable algorithms,” Ph.D. dissertation, Princeton University, 2015.
- [21] J. Frankle, S. Park, D. Shaar, S. Goldwasser, and D. Weitzner, “Practical accountability of secret processes,” in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 657–674.
- [22] G. Proudler, L. Chen, and C. Dalton, “Trusted platform architecture,” in *Trusted Computing Platforms*. Springer, 2014, pp. 109–129.
- [23] S. Johnson and D. Rizzo, “Titan silicon root of trust for google cloud,” in *Secure Enclaves Workshop*, 2018, https://keystone-enclave.org/workshop-website-2018/slides/Scott_Google_Titan.pdf.
- [24] “OpenTitan: an open source, transparent, high-quality reference design and integration guidelines for silicon root of trust (rot) chips,” <https://opentitan.org/>.
- [25] B. H. Møller, J. G. Søndergaard, K. S. Jensen, M. W. Pedersen, T. W. Bøgedal, A. Christensen, D. B. Poulsen, K. G. Larsen, R. R. Hansen, T. R. Jensen *et al.*, “Preliminary security analysis, formalisation, and verification of opentitan secure boot code,” in *Nordic Conference on Secure IT Systems*. Springer, 2021, pp. 192–211.
- [26] “PAD Places - a location sharing app on ios and google implementing privacy-preserving accountable decryption,” <https://www.pad.tech/pad-places>, 2022.
- [27] J. Parra Moyano and O. Ross, “KYC optimization using distributed ledger technology,” *Business & Information Systems Engineering*, vol. 59, no. 6, pp. 411–423, 2017.
- [28] 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 (SP)*. IEEE, 2017, pp. 483–502.
- [29] N. Kobeissi, K. Bhargavan, and B. Blanchet, “Automated verification for secure messaging protocols and their implementations: A symbolic and computational approach,” in *2017 IEEE European symposium on security and privacy (EuroS&P)*. IEEE, 2017, pp. 435–450.
- [30] N. Kobeissi, G. Nicolas, and K. Bhargavan, “Noise explorer: Fully automated modeling and verification for arbitrary noise protocols,” in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 356–370.

- [31] B. Blanchet, "Symbolic and computational mechanized verification of the arinc823 avionic protocols," in *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. IEEE, 2017, pp. 68–82.
- [32] V. Cortier, D. Galindo, and M. Turuani, "A formal analysis of the neuchâtel e-voting protocol," in *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2018, pp. 430–442.
- [33] B. Blanchet, B. Smyth, V. Cheval, and M. Sylvestre, "Proverif 2.04: automatic cryptographic protocol verifier, user manual and tutorial," *Version from*, 2021.
- [34] M. Abadi and V. Cortier, "Deciding knowledge in security protocols under equational theories," *Theoretical Computer Science*, vol. 367, no. 1, pp. 2–32, 2006.
- [35] V. Cheval, V. Cortier, and M. Turuani, "A little more conversation, a little less action, a lot more satisfaction: Global states in ProVerif," in *IEEE Computer Security Foundations Symposium (CSF)*, Jul. 2018, pp. 344–358.
- [36] 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: ACM Press, Nov. 2022.
- [37] B. Blanchet, V. Cheval, and V. Cortier, "Proverif with lemmas, induction, fast subsumption, and much more (long version)," <https://chevalvi.gitlabpages.inria.fr/chevalvi/files/BCC-snp22-long.pdf>, Inria Paris, CNRS, LORIA, Tech. Rep., 2021.

Appendix A.

Detailed syntax & semantics

We reuse most of the notation given in [37]. We simplify however the definitions as we do not consider here equivalence properties, phases or tables. Recall that \mathcal{C}_{user} is the set of Horn clauses given in the input file describing the semantics of the predicates in \mathcal{F}_p and that for all $p \in \mathcal{F}_p$, the semantics of p , i.e. $\text{sem}(p)$, is defined by the derivability within \mathcal{C}_{user} . Additionally, we consider a set of user-given abstract predicates (i.e. predicates declared with the option “blocking”) \mathcal{F}_{ap} such that for all $p \in \mathcal{F}_{ap}$, the set of $\text{sem}(p)$ is assumed to exist. As mentioned in the body of the paper, we will consider the set \mathcal{F}_{bp} of b- p with $p \in \mathcal{F}_p$. Note that given $p \in \mathcal{F}_{ap}$, we define $\text{b-}p = p$ (the predicate is already blocking).

A.1. Semantics

Let \mathcal{V} be an infinite set of variables and \mathcal{N} be an infinite set of names. The semantics of a destructor function symbol g is given by a sequence of rewrite rules $\text{def}(g) = [g(U_{i,1}, \dots, U_{i,n}) \rightarrow U_i]_{i=1}^k$ where $U_{i,j}$ and U_i can be either terms or the special constant fail. The result of the evaluation of a term $g(M_1, \dots, M_n)$ is typically the result of the first rewrite rule in $\text{def}(g)$ that matches the arguments M_1, \dots, M_n . Formally, the evaluation of an expression D to U , denoted $D \Downarrow U$, is given by:

- if $D \in \mathcal{V} \cup \mathcal{N} \cup \{\text{fail}\}$ then $D \Downarrow D$;
- if $D = f(D_1, \dots, D_n)$ and $D_1 \Downarrow V_1, \dots, D_n \Downarrow V_n$:
 - if $f \in \mathcal{F}_c$ and $\text{fail} \in \{V_1, \dots, V_n\}$ then $D \Downarrow \text{fail}$;
 - otherwise if $f \in \mathcal{F}_c$ then $D \Downarrow f(V_1, \dots, V_n)$;
 - otherwise $f \in \mathcal{F}_d$ and we let $D' = g(V_1, \dots, V_n)$. Then either D' cannot be rewritten by any rule of $\text{def}(g)$ and $D \Downarrow \text{fail}$, or otherwise $D \Downarrow V$ where $D' \rightarrow V$ by the first rule of $\text{def}(g)$ applicable to D' .

The semantics of processes is given in Fig. 7.

A.2. Satisfaction of facts and queries

To express secrecy and correspondence properties, ProVerif consider facts as follows:

$$\begin{array}{ll}
 F ::= & \text{event}(ev(M_1, \dots, M_n)) \quad ev \text{ event} \\
 & \text{mess}(M, N) \\
 & \text{att}(M) \\
 & p(M_1, \dots, M_n) \quad p \in \mathcal{F}_p \\
 & M = N \\
 & M \neq N
 \end{array}$$

The satisfaction relation \vdash is defined over a trace T , a step τ and a ground fact F , denoted $T, \tau \vdash F$, as follows:

- $T, \tau \vdash \text{att}(M)$ iff $T[\tau] \rightarrow^* \mathcal{E}, \mathcal{P}, \mathcal{A}$ by only the rules APP and NEW such that $M \in \mathcal{A}$.
- $T, \tau \vdash \text{mess}(M, N)$ iff $T[\tau - 1] \xrightarrow{\text{mess}(M, N)} T[\tau]$.
- $T, \tau \vdash \text{event}(ev)$ iff $T[\tau - 1] \xrightarrow{\text{event}(ev)} T[\tau]$.
- $T, \tau \vdash p(M_1, \dots, M_n)$ iff $p(M_1, \dots, M_n) \in \text{sem}(p)$.with $p \in \mathcal{F}_p \cup \mathcal{F}_{ap}$.
- $T, \tau \vdash M = N$ iff $M = N$.
- $T, \tau \vdash M \neq N$ iff $M \neq N$.

Note that only facts $\text{mess}(M, N)$, $\text{event}(ev)$ and $\text{att}(M)$ actually depend on the step τ and the trace T . Thus, we

also write $\vdash p(M_1, \dots, M_n)$, $\vdash M = N$ and $\vdash M \neq N$. Finally, when the exact value of the step is unnecessary, we can write $T \vdash F$ instead of $\exists \tau. T, \tau \vdash F$. The satisfaction of first order logic formula over facts is naturally defined based on \vdash .

As mentioned in the body of the paper, by seen a correspondence query $F_1 \wedge \dots \wedge F_n \Rightarrow \psi$ as a first order logic formula $\Psi = \forall \tilde{x}. (F_1 \wedge \dots \wedge F_n \Rightarrow \exists \tilde{y}. \psi)$ with $\tilde{x} = \text{vars}(F_1, \dots, F_n)$ and $\tilde{y} = \text{vars}(\psi) \setminus \tilde{x}$, its semantics is defined as: for all traces T of P , $T \vdash \Psi$.

A.3. From operational to instrumented semantics

Freshness of names from rules NEW and RESTR have always been problematic to handle when it comes to automatic verification. In [37], they deal with the problem by *instrumenting* their syntax and semantics. Typically, replication processes $!P$ will be associated with a session variable i , i.e. $!^i P$, and processes $\text{new } k; P$ will be replaced by $P\{k[\tilde{x}]/k\}$ where $k[\tilde{x}]$ represents a pattern and \tilde{x} are the input and session variables in the scope of $\text{new } k$. In other words, the name becomes a function that is applied on its environment at the time its creation. Such instrumentation is shown to be correct [37, Lemma 3] where the satisfaction relation on instrumented facts is denoted \vdash_i . We naturally extend the satisfaction relation to user-defined predicates as follows: For all $p \in \mathcal{F}_p \cup \mathcal{F}_{ap}$, for all T, τ , $T, \tau \vdash_i p(M_1, \dots, M_n)$ iff $T, \tau \vdash_i \text{b-}p(M_1, \dots, M_n)$ iff $p(M_1, \dots, M_n) \in \text{sem}(p)$.

It is in fact shown that we can also restrict the space of traces when verifying a correspondence query, that is by only considering IO compliant traces [37, Definition 13]. The exact details of the definition is not relevant to this paper but it intuitively restricts the satisfaction of attacker facts to what is exactly known at a current step (and not what the attacker can learn from its knowledge). Similarly, it forces the satisfaction of message fact to go through the attacker and not through an internal communication when the query does not request to prove message facts. Note that the IO compliance of a trace and the IO compliance of a query is normally defined w.r.t. a phase but as previously mentioned, we simplified the framework by removing tables and phases.

We consider the same satisfaction relation \vdash_{IO} that we extend to user-defined predicates as follows:

- $T, \tau \vdash_{IO} p(M_1, \dots, M_n)$ iff $p(M_1, \dots, M_n) \in \text{sem}(p)$, $\tau = \min_{\mathcal{C}_{user} \cup \mathcal{C}_{bp}}(p(M_1, \dots, M_n))$, and $p \in \mathcal{F}_p$
- $T, \tau \vdash_{IO} \text{b-}p(M_1, \dots, M_n)$ with $p \in \mathcal{F}_p \cup \mathcal{F}_{ap}$ iff $\tau = 0$ and $p(M_1, \dots, M_n) \in \text{sem}(p)$.

where $\mathcal{C}_{bp} = \{\rightarrow F \mid \text{pred}(F) = p \in \mathcal{F}_{bp} \cup \mathcal{F}_{ap} \wedge F \in \text{sem}(p)\}$ and $\min_{\mathcal{C}_{user} \cup \mathcal{C}_{bp}}(p(M_1, \dots, M_n))$ is the size (in term of number of nodes) of the smallest derivation of $p(M_1, \dots, M_n)$ in $\mathcal{C}_{user} \cup \mathcal{C}_{bp}$.

Note that contrary to the satisfaction relation \vdash , in $T, \tau \vdash_{IO} p(M_1, \dots, M_n)$, the integer τ provides some information on the deduced predicate, namely that it corresponds to the size of the smallest derivation capable of deriving the fact. Thus, τ does not have the *step* signification as it is for attacker, message and event facts. Note that the satisfaction is still independent from the trace T .

$\mathcal{E}, \mathcal{P} \cup \{0\}, \mathcal{A} \rightarrow \mathcal{E}, \mathcal{P}, \mathcal{A}$	(NIL)
$\mathcal{E}, \mathcal{P} \cup \{P \mid Q\}, \mathcal{A} \rightarrow \mathcal{E}, \mathcal{P} \cup \{P, Q\}, \mathcal{A}$	(PAR)
$\mathcal{E}, \mathcal{P} \cup \{!P\}, \mathcal{A} \rightarrow \mathcal{E}, \mathcal{P} \cup \{P, !P\}, \mathcal{A}$	(REPL)
$\mathcal{E}, \mathcal{P} \cup \{\text{new } a; P\}, \mathcal{A} \rightarrow \mathcal{E} \cup \{a'\}, \mathcal{P} \cup \{P\{a'/a\}\}, \mathcal{A}$	if $a' \notin \mathcal{E}$ (RESTR)
$\mathcal{E}, \mathcal{P} \cup \{\text{out}(N, M); P, \text{in}(N, x); Q\}, \mathcal{A} \xrightarrow{\text{mess}(N, M)} \mathcal{E}, \mathcal{P} \cup \{P, Q\{M/x\}\}, \mathcal{A}$	(I/O)
$\mathcal{E}, \mathcal{P}, \mathcal{A} \xrightarrow{\text{mess}(N, M)} \mathcal{E}, \mathcal{P}, \mathcal{A}$	if $N, M \in \mathcal{A}$ (MSG)
$\mathcal{E}, \mathcal{P} \cup \{\text{let } x = D \text{ in } P \text{ else } Q\}, \mathcal{A} \rightarrow \mathcal{E}, \mathcal{P} \cup \{P\{M/x\}\}, \mathcal{A}$	if $D \Downarrow M$ (LET1)
$\mathcal{E}, \mathcal{P} \cup \{\text{let } x = D \text{ in } P \text{ else } Q\}, \mathcal{A} \rightarrow \mathcal{E}, \mathcal{P} \cup \{Q\}, \mathcal{A}$	if $D \Downarrow \text{fail}$ (LET2)
$\mathcal{E}, \mathcal{P} \cup \{\text{out}(N, M); P\}, \mathcal{A} \xrightarrow{\text{mess}(N, M)} \mathcal{E}, \mathcal{P} \cup \{P\}, \mathcal{A} \cup \{M\}$	if $N \in \mathcal{A}$ (OUT)
$\mathcal{E}, \mathcal{P} \cup \{\text{in}(N, x); Q\}, \mathcal{A} \xrightarrow{\text{mess}(N, M)} \mathcal{E}, \mathcal{P} \cup \{Q\{M/x\}\}, \mathcal{A}$	if $N, M \in \mathcal{A}$ (IN)
$\mathcal{E}, \mathcal{P}, \mathcal{A} \rightarrow \mathcal{E}, \mathcal{P}, \mathcal{A} \cup \{M\}$	(APP)
$\mathcal{E}, \mathcal{P}, \mathcal{A} \rightarrow \mathcal{E} \cup \{a'\}, \mathcal{P}, \mathcal{A} \cup \{a'\}$	if $a' \notin \mathcal{E}$ (NEW)
$\mathcal{E}, \mathcal{P} \cup \{\text{event}(ev); P\}, \mathcal{A} \xrightarrow{\text{event}(ev)} \mathcal{E}, \mathcal{P} \cup \{P\}, \mathcal{A}$	(EVENT)
$\mathcal{E}, \mathcal{P} \cup \{\text{let } x_1, \dots, x_n \text{ suchthat } p(M_1, \dots, M_k) \text{ in } P \text{ else } Q\} \rightarrow \mathcal{E}, \mathcal{P} \cup \{P\sigma\}, \mathcal{A}$	(PRED1)
$\mathcal{E}, \mathcal{P} \cup \{\text{let } x_1, \dots, x_n \text{ suchthat } p(M_1, \dots, M_k) \text{ in } P \text{ else } Q\} \rightarrow \mathcal{E}, \mathcal{P} \cup \{Q\}, \mathcal{A}$	(PRED2)
$\text{if } \text{dom}(\sigma) = \{x_1, \dots, x_n\} \text{ and } p(M_1\sigma, \dots, M_k\sigma) \in \text{sem}(p)$	
$\text{if for all } \sigma \text{ with } \text{dom}(\sigma) = \{x_1, \dots, x_n\}, p(M_1\sigma, \dots, M_k\sigma) \notin \text{sem}(p)$	

Figure 7. Transitions between configurations.

The satisfaction of a correspondence query is the same as in [37, Definition 10]. We recall it in the simplified version where we do not consider injective events nor nested queries, as they are not the main goal of this paper.

Definition 2. Given two satisfaction relations \vdash_p and \vdash_c , given a set of traces \mathcal{T} of some initial instrumented configuration \mathcal{C}_I , we say that \mathcal{T} satisfies ϱ w.r.t. to \vdash_p and \vdash_c , denoted $(\vdash_p, \vdash_c, \mathcal{T})$, when:

For all traces $T \in \mathcal{T}$, for all tuples of integers $\tilde{\tau} = (\tau_1, \dots, \tau_n)$, for all substitutions σ , if $T, \tau_i \vdash_p F_i\sigma$ for $i = 1, \dots, n$ then there exists a substitution σ' such that $F_i\sigma = F_i\sigma'$ for $i = 1, \dots, n$ and $T \vdash_c \psi\sigma'$.

We distinguish two satisfaction relations, one for the premises of the query, i.e. \vdash_p , and the other for the conclusion of the query, i.e. \vdash_c . This separation originally occurred in [37] as a proof technique. But in this case, we will in fact use to show the different way we need to prove lemmas for them to be used correctly.

We obtain the same result linking the restrictions to the main satisfiability relation:

Lemma 1 ([37, Lemma 8]). Let $\mathcal{C}_I = \rho, P, \mathcal{A}$ be an initial instrumented configuration. Let ϱ be an IO-compliant correspondence query such that $\text{names}(\varrho) \subseteq \text{dom}(\rho)$. We have:

$$\begin{aligned} (\vdash_i, \vdash_i, \text{trace}(\mathcal{C}_I)) \models \varrho \\ \text{if and only if} \\ (\vdash_{IO}, \vdash_i, \text{trace}_{IO}(\mathcal{C}_I)) \models \varrho \end{aligned}$$

A.4. Proof by induction

In [37], a query is proved by induction typically on the size of the trace and the multiset of steps (τ_1, \dots, τ_n) from Definition 2. Moreover, as the proof by induction

typically consists on generating a special *inductive lemma*, representing the inductive hypothesis, which would be applied during the saturation procedure, ProVerif currently simplifies the conclusion of the query to remove non-blocking events when generating such inductive lemma. In our case, thanks to our generalization of blocking predicate, we do not need to simplify the conclusion query but we will replace all predicates in the conclusion of the query with their blocking counterpart. However, our inductive hypothesis had to be extended as stated below.

Definition 3. Let ϱ be a IO-compliant query of the form $\bigwedge_{i=1}^n F_i \Rightarrow \psi$. For all traces T , for all tuples of steps $\tilde{\tau} = (\tau_1, \dots, \tau_k)$ and $\tilde{\tau}_p = (\tau_{k+1}, \dots, \tau_m)$, $\mathcal{I}\mathcal{H}_\varrho(T, \tilde{\tau}, \tilde{\tau}_p)$ holds, if and only if, if $n = m$ and for all $j \leq n$, $\text{pred}(F_j) \in \mathcal{F}_p \Leftrightarrow j \geq k + 1$ then

for all substitutions σ , if $T, \tau_i \vdash_{IO} F_i\sigma$ for $i = 1, \dots, n$ then there exists σ' such that $F_i\sigma = F_i\sigma'$ for $i = 1, \dots, n$ and $T \vdash_{IO} [\psi\sigma']^b$.

There are two key differences with [37, Definition 15]: (i) we separated the steps from standard facts, att, mess and event, with the ones corresponding to predicate from \mathcal{F}_p . (ii) the satisfaction of the conclusion of the query relies on \vdash_{IO} instead of \vdash_i , the main difference being related to the satisfaction of $\text{att}(M)$ facts. Intuitively, $T, \tau \vdash_i \text{att}(M)$ holds if the attacker can create M from its knowledge at τ , whereas $T, \tau \vdash_{IO} \text{att}(M)$ ensures that the attacker already has M in its knowledge. This distinction matter for the proof of induction has we need to order the facts, so that we can apply our inductive hypothesis on them.

To state the properties for lemmas and inductive lemmas needed for expressing the soundness of the saturation, we naturally extended the strict order relation $<_{ind}$ to

the lexicographic ordering of triplets $(T, \tilde{\tau}, \tilde{\tau}_p)$ where comparison of the size of T is used for the first element, and multiset ordering on natural number is used for the second and third element, i.e. $\tilde{\tau}$ and $\tilde{\tau}_p$ are viewed as a multiset of natural number instead of tuples in the order.

The properties on lemmas and inductive lemmas needed for the soundness can thus be summarized as follow.

Definition 4 ([37, Definition 22]). *Let \mathcal{C}_I be an initial instrumented configuration. Let $\mathcal{L}, \mathcal{L}_i$ be two sets of lemmas. Let $\tilde{\tau}$ and $\tilde{\tau}_p$ be two tuples of steps. Let $T \in \text{trace}_{IO}(\mathcal{C}_I)$. We define the predicate $\text{Hyp}_{\mathcal{L}, \mathcal{L}_i}(T, \tilde{\tau}, \tilde{\tau}_p)$ to hold if and only if*

- for all $\varrho \in \mathcal{L}$, $(\vdash_{IO}, \vdash_{IO}, \text{trace}_{IO}(\mathcal{C}_I)) \models \varrho$;
- for all $\varrho \in \mathcal{L}_i$, for all $T' \in \text{trace}_{IO}(\mathcal{C}_I)$, for all tuples of steps $\tilde{\tau}', \tilde{\tau}'_p$, if $(T', \tilde{\tau}', \tilde{\tau}'_p) <_{\text{ind}} (T, \tilde{\tau}, \tilde{\tau}_p)$ then $\mathcal{IH}_{\varrho}(T', \tilde{\tau}', \tilde{\tau}'_p)$.

As in [37], the set \mathcal{L} in the above definition represents the lemmas already proved by ProVerif (or the axioms) whereas \mathcal{L}_i are the query that we are currently trying to prove by induction.

Appendix B. The procedure

B.1. Generation of Horn clauses

The generation of Horn clauses is left unchanged from [37] and are defined by the translation $\llbracket P, \mathcal{I} \rrbracket \mathcal{H}\rho$ where P is a process, \mathcal{I} is a sequence of session variables and terms, \mathcal{H} is a conjunction of facts and disequalities for the form $\forall \tilde{x}. M \neq N$, and ρ is a mapping from names to instrumented names. We present in Fig. 8 the translation fitting our simplified framework.

$$\begin{aligned}
\llbracket 0, \mathcal{I} \rrbracket \mathcal{H}\rho &= \emptyset \\
\llbracket P \mid Q, \mathcal{I} \rrbracket \mathcal{H}\rho &= \llbracket P, \mathcal{I} \rrbracket \mathcal{H}\rho \cup \llbracket Q, \mathcal{I} \rrbracket \mathcal{H}\rho \\
\llbracket !P, \mathcal{I} \rrbracket \mathcal{H}\rho &= \llbracket P, (\mathcal{I}, i) \rrbracket \mathcal{H}\rho \\
&\text{where } i \text{ is a fresh variable session identifier} \\
\llbracket \text{new } a; P, \mathcal{I} \rrbracket \mathcal{H}\rho &= \llbracket P, \mathcal{I} \rrbracket \mathcal{H}(\rho[a \mapsto a[\mathcal{I}]]) \\
\llbracket \text{in}(M, x); P, \mathcal{I} \rrbracket \mathcal{H}\rho &= \\
&\llbracket P, (\mathcal{I}, x') \rrbracket (\mathcal{H} \wedge \text{mess}(M\rho, x'))(\rho[x \mapsto x']) \\
\llbracket \text{out}(M, N); P, \mathcal{I} \rrbracket \mathcal{H}\rho &= \\
&\llbracket P, \mathcal{I} \rrbracket \mathcal{H}\rho \cup \{\mathcal{H} \rightarrow \text{mess}(M\rho, N\rho)\} \\
\llbracket \text{let } x = D \text{ in } P \text{ else } Q, \mathcal{I} \rrbracket \mathcal{H}\rho &= \\
&\bigcup \{ \llbracket P, \mathcal{I}\sigma \rrbracket (\mathcal{H}\sigma \wedge \phi)(\rho\sigma[x \mapsto M]) \mid D\rho \Downarrow' (M, \sigma, \phi) \} \\
&\cup \bigcup \{ \llbracket Q, \mathcal{I}\sigma \rrbracket (\mathcal{H}\sigma \wedge \phi)(\rho\sigma) \mid D\rho \Downarrow' (\text{fail}, \sigma, \phi) \} \\
\llbracket \text{event}(ev); P, \mathcal{I} \rrbracket \mathcal{H}\rho &= \\
&\llbracket P, \mathcal{I} \rrbracket (\mathcal{H} \wedge \text{s-event}(ev\rho)) \rho \cup \{\mathcal{H} \rightarrow \text{m-event}(ev\rho)\} \\
\llbracket \text{let } x_1, \dots, x_n \text{ suchthat } \text{pred} \text{ in } P \text{ else } Q, \mathcal{I} \rrbracket \mathcal{H}\rho &= \\
&\llbracket P, \mathcal{I} \rrbracket (\mathcal{H} \wedge \text{pred}\rho')(\rho\rho') \cup \llbracket Q, \mathcal{I} \rrbracket \mathcal{H}\rho \\
&\text{where } \rho' = [x_i \mapsto x'_i]_{i=1}^n \text{ and } x'_i \text{ fresh variables.}
\end{aligned}$$

Figure 8. Generation of clauses

Given an initial instrumented configuration \mathcal{C}_I , the set of clauses generated from it is denoted $\mathbb{C}_{\mathcal{P}}(\mathcal{C}_I)$. We also

give back the set of clauses generated for the attacker, denoted $\mathbb{C}_{\mathcal{A}}(\mathcal{C}_I)$, as follows:

$$\begin{aligned}
&\text{For all } a \in \mathcal{A}(\mathcal{C}_I), \rightarrow \text{att}(a[]) && \text{(RInit)} \\
&\text{att}(b_0[i]) \quad \text{with } b_0 \text{ not occurring in } \mathcal{C}_I && \text{(RGen)} \\
&\text{att}(\text{fail}) && \text{(RFail)} \\
&\text{For all functions } h, \\
&\quad \text{for all } h(U_1, \dots, U_n) \rightarrow U \mid \phi \text{ in } \text{def}(h) && \text{(Rf)} \\
&\quad \text{att}(U_1) \wedge \dots \wedge \text{att}(U_n) \wedge \phi \rightarrow \text{att}(U) \\
&\text{mess}(x, y) \wedge \text{att}(x) \rightarrow \text{att}(y) && \text{(RI)} \\
&\text{att}(x) \wedge \text{att}(y) \rightarrow \text{mess}(x, y) && \text{(Rs)}
\end{aligned}$$

Given a trace T of an initial instrumented configuration \mathcal{C}_I , we denote by $\mathbb{C}_b(T)$ the set of Horn clauses of the form $\rightarrow \text{b-p}(M_1, \dots, M_n)$ such that $T, \tau \vdash_{IO} \text{b-p}(M_1, \dots, M_n)$ for some step τ . Note that $\mathbb{C}_b(T)$ contains also \mathbb{C}_{b_p} .

B.2. Derivation and initial soundness

We preserve exactly the notion of derivation of Horn clauses used in the saturation but the satisfaction relation of a derivation by a trace had to be extended.

Definition 5 ([37, Definition 17]). *Let \mathbb{C} be a set of clauses. Let F be a closed fact and a step τ . A derivation \mathcal{D} of F at step τ from \mathbb{C} is a finite tree defined as follows:*

- its nodes (except the root) are labelled by clauses from \mathbb{C} .
- its edges are labelled by ground facts and by a step.
- if the tree contains a node labelled by R with one incoming edge labelled by F_0 and n outgoing edges labelled by F_1, \dots, F_n then $R \sqsupseteq F_1 \wedge \dots \wedge F_n \rightarrow F_0$.
- the root has one outgoing edge, labelled by F and τ .

In the above definition, $R \sqsupseteq R'$ denotes the multiset subsumption of R' by R .

We define the *blocking-free size* of a derivation \mathcal{D} as the number of nodes in \mathcal{D} not labelled by a clause of the form $\rightarrow C$ with $\text{pred}(C) = \text{b-p}$ for some p . The notion of satisfaction of a derivation by a trace is also parametrized by a set of predicate \mathcal{S}_p on which we need to have guarantees on their satisfiability in the trace w.r.t. \vdash_{IO} . Typically, \mathcal{S}_p will contain the predicate att , mess and $p \in \mathcal{F}_p$ if they occur in the conclusion of the query or in the premise of a lemma. It will also contain all the blocking predicates and the user-defined predicate \mathcal{F}_{ap} .

Definition 6. *Let \mathcal{D} be a derivation of F at step τ . Let \mathcal{S}_p be a set of predicates. We say that a trace T satisfies a derivation \mathcal{D} w.r.t. \mathcal{S}_p , denoted $T, \mathcal{S}_p \vdash \mathcal{D}$, when for all nodes η of \mathcal{D} , if F_0, τ_0 is the label of the incoming edge of η then*

- 1) if η is labelled with **(RI)** then the outgoing edges of η are labelled $\text{mess}(N, M), \tau'$ and $\text{att}(N), \tau''$ for some N, M, τ', τ'' such that $\tau' \leq \tau_0$ and if $\text{att} \in \mathcal{S}_p$ then $\tau'' < \tau_0$ else $\tau'' \leq \tau_0$.
- 2) if η is not labelled with **(RI)**, then for all outgoing edges of η labelled F', τ' , if $\text{pred}(F_0) \in \mathcal{F}_p$ or $\text{pred}(F') \notin \mathcal{F}_p$ then

- $\text{pred}(F') \notin \mathcal{S}_p$ implies $\tau' \leq \tau_0$
- $\text{pred}(F') \in \mathcal{S}_p$ implies $\tau' < \tau_0$

- 3) if $F_0 = \text{att}(f(M_1, \dots, M_m))$ with $f \in \mathcal{F}_{data}$ then
- either R is the clause $\text{att}(x_1) \wedge \dots \wedge \text{att}(x_m) \rightarrow \text{att}(f(x_1, \dots, x_m))$ and if $\text{att} \in \mathcal{S}_p$ then $\tau_0 = \min\{\tau \mid T, \tau \vdash_{IO} F_0\}$;
 - or η is not the root and the node η' connected to the incoming edge of η is labelled with the clause $\text{att}(f(x_1, \dots, x_m)) \rightarrow \text{att}(x_j)$ for some j and if $\text{att} \in \mathcal{S}_p$ then $f(M_1, \dots, M_m) \in \mathcal{A}(T[\tau_0])$;
- else $\text{pred}(F_0) \in \mathcal{S}_p$ implies $\tau_0 = \min\{\tau \mid T, \tau \vdash_{IO} F_0\}$

The main different before our definition and [37, Definition 18] lies in Item 2 where we relaxed the ordering conditions between incoming and outgoing edges, but only in the particular case where the incoming edge is labelled by standard predicates, att or mess, and the outgoing edges is labelled by a predicate fact F with $\text{pred}(F)$ being a user-defined predicate. The clause labeling such a node typically only comes from the translation of the process into Horn clauses (recall that \mathbb{C}_{user} cannot contain att nor mess). However, in the satisfaction relation $T, \tau \vdash_{IO} F$, the step τ represents an actual step of the trace T when $\text{pred}(F)$ are standard predicates whereas τ represents the size of the smallest derivation in \mathbb{C}_{user} when $\text{pred}(F) \in \mathcal{F}_p$. These two notions being unrelated, we cannot hope to order them. When $\text{pred}(F) \in \mathcal{F}_{ap} \cup \mathcal{F}_{bp}$, the satisfaction $T, \tau \vdash_{IO} F$ implies $\tau = 0$ and so we are naturally ordered.

Note, on the order hand, when both incoming and outgoing edge are labelled by user-defined predicates (i.e. the Horn clause labeling the node originated from \mathbb{C}_{user}), we do require that the step τ_0 of the incoming edge is strictly bigger than the one of the outgoing edge (when required by \mathcal{S}_p). This will typically allow us to apply inductive hypotheses from Definition 4.

Another difference comes from Item 3 where we require τ_0 to be the minimum step that satisfies F_0 . We can basically always take such minimum by replacing the corresponding sub derivation with one proving the same fact with a smaller step. It is however useful when relating standard predicates and their blocking counterpart.

With this, we can state the first soundness result linking the satisfaction of facts in a trace and the satisfaction of a derivation. The following theorem is an adaption of [37, Theorem 1].

Theorem 2. *Let \mathcal{C}_I be an initial instrumented configuration. Let \mathcal{S}_p be a set of predicates.*

For all $T \in \text{trace}_{IO}(\mathcal{C}_I)$, for all ground facts F , for all steps τ , if $T, \tau \vdash_{IO} F$ then there exists a derivation \mathcal{D} of F at step τ from $\mathbb{C}_{\mathcal{A}}(\mathcal{C}_I) \cup \mathbb{C}_{\mathcal{P}}(\mathcal{C}_I) \cup \mathbb{C}_{\mathcal{B}}(T)$ such that $T, \mathcal{S}_p \vdash \mathcal{D}$.

Proof sketch. The proof is typically the same as in [37, Theorem 1] as we did not modify the core generation of Horn clauses and satisfiability relations but only extended them user-defined predicates. Hence we only need to discuss the new cases arose by them.

Notice that if $\text{pred}(F)$ is a blocking predicate b- p , $T, \tau \vdash_{IO} F$ directly gives us that $\rightarrow F$ is in $\mathbb{C}_{\mathcal{B}}(T)$ which allows us to conclude.

In the case where the protocol executes the construct let x_1, \dots, x_n suchthat $p(M_1, \dots, M_n)$ in P else Q , the semantics rule (PRED1) ensures that $p(M_1\sigma, \dots, M_n\sigma)$

is derivable from $\mathbb{C}_{user} \cup \mathbb{C}_{bp}$. Thus to satisfy Item 2 of Definition 6, we only need to take the smallest derivation \mathcal{D} of $p(M_1, \dots, M_n)\sigma$ from $\mathbb{C}_{user} \cup \mathbb{C}_{bp}$.

Notice that for the semantics rule (LETP2) requires for the predicate to not hold; however in the translation of the construct to Horn clauses, this condition is "lost". It is one of the strong abstraction of ProVerif: it behaves as if the else branch of a let ... suchthat construct can always be executed. This is due to the fact that ProVerif does not reason with the negation of predicates.

When $\text{pred}(F) \in \mathcal{F}_p$, we once again know that $F \in \text{sem}(p)$ which by definition implies that F is derivable from $\mathbb{C}_{user} \cup \mathbb{C}_{bp}$. By taking the smallest derivation, we can conclude. \square

B.3. Saturation procedure and its soundness

As previously mentioned in the body of this paper, the selection function is extended to prevent the selection of any blocking predicate facts (including the ones from \mathcal{F}_{ap}) and attacker facts $\text{att}(x)$ with x a variable.

We already gave some details on how to apply lemma in our new settings but we provide the complete definition here, as well as the application of inductive lemmas and some updated classic transformation rules.

Lemmas. Let us denote \mathbb{C}_{std} the set of clauses containing the clauses (RI), (Rf _{g}) and (Rf _{π_i^g}) for all $g \in \mathcal{F}_{data}$ and i . The following transformation rule (Lem($\mathcal{L}, \mathcal{S}_p$)) describes the application of (non-inductive) lemmas in \mathcal{L} with the set of predicate \mathcal{S}_p .

$$\frac{\mathbb{C} \cup \{R = (H \rightarrow C)\} \quad R \notin \mathbb{C}_{std} \quad \begin{array}{l} (\bigwedge_{i=1}^n F_i \Rightarrow \bigvee_{j=1}^m \psi_j) \in \mathcal{L} \\ \forall i, \text{pred}(F_i) \in \mathcal{S}_p \text{ and either } [F_i\sigma]^b \in [H]^b \\ \text{or } (F_i\sigma = C \text{ and } \forall j, \forall F \in \psi_j, \text{mgu}(F\sigma, C) = \perp \\ \text{or } (F_i\sigma = C \text{ and } \text{pred}(C) \in \mathcal{F}_p) \end{array}}{\mathbb{C} \cup \{H \wedge [\psi_j\sigma]^b \rightarrow C\}_{j=1}^m}$$

Proved lemmas can be applied on both hypotheses and the conclusion. However, when applied thanks to the conclusion, we want to preserve the invariant that the facts in the hypotheses of a Horn clause occurs strictly before the conclusion of the clause. However, a lemma only guarantees us that the facts in $\bigvee_{j=1}^m \psi_j$ occur before or at the same time as F_i . Hence the condition on unification of facts from the conclusion of the query with the conclusion of the clause. Note that this does not apply when the conclusion of the clause is rooted by a user-defined predicate as a lemma only adds blocking predicates from $\mathcal{F}_{bp} \cup \mathcal{F}_{ap}$ which by definition of the satisfaction relation \vdash_{IO} are always satisfied at step 0.

Inductive lemmas. The application inductive lemmas differ from normal lemmas by the fact that we do not try to match the conclusion of the Horn clause with the premise of the inductive lemma. Recall that an inductive lemma intuitively represents the inductive hypothesis. Hence we can only apply it on facts occurred strictly before the conclusion. For the same reason, we also do not match the premise of the inductive lemma with a fact rooted by a blocking predicate b- p with $p \in \mathcal{F}_p$. Indeed, when such predicate occurs in the hypothesis of a clause, it means that it was added by a previous application of a lemma

(inductive or not). However, our induction for user-defined predicate is done typically on the size of the minimal derivation in $\mathbb{C}_{user} \cup \mathbb{C}_{bp}$. Queries and lemmas do not guarantee any relation on the size of derivations between facts of their premise and their conclusion. Hence, we must not apply our inductive lemma on these facts. The following transformation rule ($\text{Ind}(\mathcal{L}_i, \mathcal{S}_p)$) describes the application of inductive lemmas from \mathcal{L}_i with the set of predicates \mathcal{S}_p .

$$\frac{\mathbb{C} \cup \{R = (H \rightarrow C)\} \quad R \notin \mathbb{C}_{std} \quad \left(\bigwedge_{i=1}^n F_i \Rightarrow \bigvee_{j=1}^m \psi_j \right) \in \mathcal{L}_i \quad \forall i, \text{pred}(F_i) \in \mathcal{S}_p \text{ and either } F_i \sigma \in H \text{ or } (\text{pred}(F_i) \notin \mathcal{F}_p \text{ and } \lceil F_i \sigma \rceil^b \in \lceil H \rceil^b)}{\mathbb{C} \cup \{H \wedge \lceil \psi_j \sigma \rceil^b \rightarrow C\}_{j=1}^m}$$

Classic transformation rule. As previously mentioned in the body of this paper, we can amend some of the classic simplification rules to benefit from the blocking facts introduced by lemmas. Note that all these extended simplifications rules are sound but it does not mean that they should always be applied. Depending on the needs (precision, termination, ...), one may prefer to rely on the usual rules.

Consider for the tautology rule:

$$\frac{\mathbb{C} \cup \{F' \wedge H \rightarrow F\} \quad F = F' \text{ or } (\text{pred}(F) \notin \mathcal{F}_p \text{ and } \lceil F \rceil^b = \lceil F' \rceil^b)}{\mathbb{C}}$$

This rule does not make the distinction between facts and their blocking counter fact when the predicate rooting F is standard, i.e. not in \mathcal{F}_p . Indeed, for these predicates, the application of a lemma preserves the ordering, i.e. we know that F' occurred strictly before F in the sense of trace steps. When F is rooted by \mathcal{F}_p however, its blocking counterpart is not ordered with respect to the conclusion and can be seen as "additional information" that can be propagated later. As such, it does not necessarily imply that such clause with $F' = \lceil F \rceil^b$ and $\text{pred}(F) \in \mathcal{F}_p$ is useless.

We update the redundant hypotheses rule as follows.

$$\frac{\mathbb{C} \cup \{H' \wedge H \wedge \phi \rightarrow C\} \quad \lceil H' \sigma \rceil^b \subseteq \lceil H \rceil^b \quad \text{dom}(\sigma) \cap \text{vars}(H, C) = \emptyset \quad \phi \models \phi \sigma}{\mathbb{C} \cup \{H \wedge \phi \sigma \rightarrow C\}}$$

We also update the attacker rule:

$$\frac{\mathbb{C} \cup \{F \wedge H \wedge \phi \rightarrow C\} \quad \lceil F \rceil^b = \mathbf{b}\text{-att}(x) \quad x \text{ does not appear in } H, C}{\mathbb{C} \cup \{H \wedge \phi \rightarrow C\}}$$

Subsumption. We can also adjust the definition of subsumption of clauses to make use of blocking predicates. We expect this modification to have a significant positive impact on the termination of ProVerif but it may also have a negative impact on its precision.

Definition 7. Let $H_1 \wedge \phi_1 \rightarrow C_1$ and $H_2 \wedge \phi_2 \rightarrow C_2$ be two clauses. We say that $H_1 \wedge \phi_1 \rightarrow C_1$ subsumes $H_2 \wedge \phi_2 \rightarrow C_2$, denoted $(H_1 \wedge \phi_1 \rightarrow C_1) \sqsupseteq_b (H_2 \wedge \phi_2 \rightarrow C_2)$, when there exists $\sigma, n, F_1, \dots, F_n, G_1, \dots, G_n$ such that

- $C_1 \sigma = C_2$
- $H_2 = G_1 \wedge \dots \wedge G_n \wedge H_2'$
- $H_1 = F_1 \wedge \dots \wedge F_n$
- for all $i \in \{1, \dots, n\}$, either $F_i \sigma = G_i$ or $F_i \sigma = \lceil G_i \rceil^b$
- $\phi_2 \models \phi_1 \sigma$.

Notice that a fact rooted by b - p can subsume a fact rooted by p . However, the converse is not true, that is a fact rooted by a non-blocking predicate p cannot subsume its blocking counterpart b - p .

Soundness of the saturation. For the rest of the saturation procedure, we proceed as in the current version of ProVerif. In particular, all other simplification rules remain unchanged (for natural number, data constructor, etc) as well as the general redundancy [37, Section 5.3].

We say that a set of clauses \mathbb{C} is consistent with \mathcal{F}_p when for all $(H \rightarrow C) \in \mathbb{C}$,

- 1) $\text{pred}(C) \in \mathcal{F}_{ap} \cup \mathcal{F}_{bp}$ implies $H = \emptyset$;
- 2) $\text{pred}(C) \in \mathcal{F}_p$ implies for all $F \in H$, $\text{pred}(F) \in \mathcal{F}_p \cup \mathcal{F}_{ap} \cup \mathcal{F}_{bp}$.

This property is directly satisfied by \mathbb{C}_{user} and the initial clauses generated from the protocol. It is also preserved by the saturation procedure.

Theorem 3. Let \mathcal{C}_I be an initial instrumented configuration. Let \mathcal{S}_p be a set of predicates. Let $\mathcal{L}, \mathcal{L}_i$ be two sets of lemmas. Let $T \in \text{trace}_{IO}(\mathcal{C}_I)$.

For all set of clauses \mathbb{C} consistent with \mathcal{F}_p containing \mathbb{C}_{std} , for all derivations \mathcal{D} of F at step τ from $\mathbb{C} \cup \mathbb{C}_b(T)$, if $T, \mathcal{S}_p \vdash \mathcal{D}$ and

- $\text{pred}(F) \in \mathcal{F}_p$ implies $\text{Hyp}_{\mathcal{L}, \mathcal{L}_i}(T, (), (\tau))$
- $\text{pred}(F) \notin \mathcal{F}_p$ implies $\text{Hyp}_{\mathcal{L}, \mathcal{L}_i}(T, (\tau), ())$

then there exists a derivation \mathcal{D}' of F at step τ from $\text{sat}_{\mathcal{L}, \mathcal{L}_i}^{\mathcal{S}_p}(\mathbb{C}) \cup \mathbb{C}_b(T)$ such that $T, \mathcal{S}_p \vdash \mathcal{D}'$.

Proof sketch. The proof follows the proof of Theorem 2 of [37]. The main idea is to first show that if the derivation \mathcal{D} of F uses a Horn clause that is transformed by one of the transformation rules, then we can build a new derivation of \mathcal{D}' of F that relies on the Horn clauses produced by the transformation rule. The main difference is of course the addition of the additional conditions related to blocking predicates but the argument necessary for the proof are similar to the current proof. For instance, for the tautology, as previously mentioned, if \mathcal{D} was using a clause $F' \wedge H \rightarrow F''$ with $\text{pred}(F'') \notin \mathcal{F}_p$ and $\lceil F'' \rceil^b = \lceil F' \rceil^b$, we would deduce from Definition 6 that there exists τ' and τ'' such that $T, \tau' \vdash_{IO} F'$ and $T, \tau'' \vdash_{IO} F''$ and $\tau' < \tau''$. Moreover, τ' and τ'' should be the minimum values that satisfy F' and F'' respectively (see Item 3). However, as $\lceil F'' \rceil^b = \lceil F' \rceil^b$, we can deduce by definition of the satisfaction relation \vdash_{IO} that $\tau' = \tau''$ which is a contradiction with $\tau' < \tau''$. Thus, \mathcal{D} cannot use the clause $F' \wedge H \rightarrow F''$ and hence we can discard the clause.

Adapting the proofs for the other transformation rules follows the same kind of arguments. \square

B.4. Solving procedure.

The verification of a correspondence query in ProVerif relies on an extended notion of Horn clauses, namely *ordered Horn clauses*, where:

- The conclusion of the Horn clause is a special fact representing a conjunction of facts. Typically, $\bigwedge_{i=1}^n F_i$ represents the fact corresponding to the conjunction $\bigwedge_{i=1}^n F_i$. We refer the reader to section 6.1 of [37] for more details.
- The hypothesis of the Horn clauses can be ordered with respect to each fact F_i of the conclusion using a partial ordering function.

Definition 8 ([37, Definition 24]). We call ordering function a *partial function* $\delta : \mathbb{N} \rightarrow \{<, \leq\}$. We call ordered facts, denoted F^δ , a fact F annotated with an ordering function δ . Finally, an ordered clause is a clause of the form $H \rightarrow \bigwedge_{i=1}^n F_i$ where each conjunct in H is an ordered fact.

The notion of derivation of a conjunction fact $\bigwedge_{i=1}^n F_i$ and its satisfiability by a trace and a tuple of steps is extended in [37, Definition 26]. These definitions remain unchanged. In fact the only part that changes in our setting is the resolution rule where we need to take into account the changes in Item 2 of Definition 6, as well as the rules for applying lemmas and inductive lemmas.

Definition 9. Let \mathcal{S}_p be a set of predicates. We say that an ordered clause $F_1^{\delta_1} \wedge \dots \wedge F_n^{\delta_n} \wedge \phi \rightarrow C$ satisfies \mathcal{S}_p when for all $i \in \{1, \dots, n\}$, if $\text{pred}(F_i) \notin \mathcal{S}_p$ then $\delta_i \notin \text{img}(\delta_i)$.

Resolution rule. The main difference between the resolution rule during the saturation procedure and the verification procedure in [37] comes from the assignment of ordering function on the hypotheses of the clauses from the saturated set. In [37], they relied on the invariant that all facts rooted by a predicate in \mathcal{S}_p occurred strictly before the conclusion. In our case, this is not always the case (as indicated in Item 2 of Definition 6) when the hypothesis is rooted by a user-defined predicate and the conclusion of the Horn clause is rooted by a standard predicate. In that case, the ordering function associated to this fact in the hypothesis should be empty.

Given an ordering function δ , we denote by $\delta^<$ the partial function with the same domain as δ and such that $\delta^<(i) = <$ for all $i \in \text{dom}(\delta)$. Moreover, given two facts F, C and an ordering function δ , we denote by $\delta_{res}(F, C, \delta)$ the ordering function defined as:

- $\delta_{res}(F, C, \delta) = \emptyset$ if $(\text{pred}(F) = \text{b-p and } p \in \mathcal{F}_p \cup \mathcal{F}_{ap})$ or $(\text{pred}(F) \in \mathcal{F}_p \text{ and } \text{pred}(C) \notin \mathcal{F}_p)$
- otherwise $\delta_{res}(F, C, \delta) = \delta^<$ if $\text{pred}(F) \in \mathcal{S}_p$
- otherwise $\delta_{res}(F, C, \delta) = \delta$

Typically, C represents the conclusion of the saturated Horn clause, F is a fact from its hypotheses and δ is the ordering function of ordered fact on which we will apply the resolution. The first bullet point covers the case where F is a user-defined block predicate as we always assign an empty ordering function on them as well as the case where F is a user-defined predicate derived from the protocol (correspond to the negation of the condition in Item 2 of Definition 6). The second and third bullet point are respectively the consequence of the second and first bullet point in Item 2 of Definition 6.

The resolution rule ($\text{Res}_o(\mathcal{S}_p)$) is therefore defined as follows:

$$\frac{F_1 \wedge \dots \wedge F_n \rightarrow C \in \mathbb{C}_{sat} \quad \sigma = \text{mgu}(F, C) \quad R = (F^\delta \wedge H_o \rightarrow C') \in \mathbb{C} \quad F \in \text{sel}(R) \quad \text{for all } i, \delta_i = \delta_{res}(F_i, C, \delta)}{F_1^{\delta_1} \sigma \wedge \dots \wedge F_n^{\delta_n} \sigma \wedge H_o \sigma \rightarrow C' \sigma}$$

Lemma simplification rule. As explained in [37], the application conditions for a lemma during the verification procedure are basically the same as in the saturation procedure. The only difference comes from the assignment of ordering functions to the facts of the conclusion of the query. In [37], such ordering function was applied to all facts since they were events. In our case however, we need to distinguish between standard predicates and user-defined predicates. As previously mentioned, the blocking counterpart of standard predicates can be ordered whereas the ones of user-defined predicates cannot. We take this into the transformation $\lceil \psi \rceil_\delta^b$. Formally, $\lceil \psi \rceil_\delta^b$ is the conjunction built from ψ where all facts F is replaced by $\lceil F \rceil_\delta^b$ when $\text{pred}(F) \notin \mathcal{F}_p$ and by $\lceil F \rceil_\emptyset^b$ otherwise (here \emptyset refer to the ordering function with empty domain).

We now define what it means for a lemma to match an ordered Horn clause: Given an ordered clause $R = (H \rightarrow \bigwedge_{i=1}^n F'_i)$ and a lemma $\varrho = (\bigwedge_{i=1}^m F_i \Rightarrow \bigvee_{j=1}^k \psi_j)$, we say that ϱ matches R with ordering function $\delta_1, \dots, \delta_m$ and substitution σ when for all $i \in \{1, \dots, m\}$, $\text{pred}(F_i) \in \mathcal{S}_p$ and

- either there exists F'' such that $F_i^{\delta_i} \in H$ and $\lceil F'' \rceil_\delta^b = \lceil F_i \sigma \rceil_\delta^b$
- or there exists $j \in \{1, \dots, k\}$ such that $F_i \sigma = F'_j$ and $\delta_i = \{j \mapsto \leq\}$.

The application rule ($\text{Lem}_o(\mathcal{L}, \mathcal{S}_p)$) is thus defined as follows:

$$\frac{\mathbb{C} \cup \{R = (H \rightarrow \bigwedge_{i=1}^n F'_i)\} \quad \varrho = (\bigwedge_{i=1}^m F_i \Rightarrow \bigvee_{j=1}^k \psi_j) \in \mathcal{L} \quad \varrho \text{ matches } R \text{ with } \delta_1, \dots, \delta_m \text{ and } \sigma \quad \text{for all } i, \text{pred}(F_i) \notin \mathcal{F}_p \text{ implies } \delta \sqsupseteq_o \delta_i}{\mathbb{C} \cup \{H \wedge \lceil \psi_k \sigma \rceil_\delta^b \rightarrow \bigwedge_{i=1}^n F'_i \lceil \sigma \rceil_{j=1}^k\}}$$

Inductive lemma simplification rule. Applying inductive lemmas consists basically to satisfy the same application conditions as for proved lemmas but we add an additional condition showing that we satisfy the order $<_{ind}$, i.e. $(T, \tilde{\tau}', \tilde{\tau}'_p) <_{ind} (T, \tilde{\tau}, \tilde{\tau}_p)$. This can be achieved by analysing the ordering functions in the hypotheses of the clauses.

Definition 10. Let $n_a \leq n_b \in \mathbb{N}$. Let $\delta_1, \dots, \delta_m$ ordering functions with $\text{dom}(\delta_i) \subseteq \{n_a, \dots, n_b\}$ for all $i \in \{1, \dots, m\}$. We say that $\delta_1, \dots, \delta_m$ are (n_a, n_b) -equal for $\{j_1, \dots, j_m\} \subseteq \{n_a, \dots, n_b\}$ when: $m \leq n_b - n_a + 1$ and

- 1) for all $k \in \{1, \dots, m\}$, $\delta_k(j_k)$ is defined;
- 2) for all $k, k' \in \{1, \dots, m\}$, $k \neq k'$ implies $j_k \neq j_{k'}$

We say that $\delta_1, \dots, \delta_m$ are (n_a, n_b) -strict for $\{j_1, \dots, j_m\} \subseteq \{n_a, \dots, n_b\}$ when $n_a \leq n_b$ and:

- either for all $k \in \{1, \dots, m\}$, there exists $i \in \{n_a, \dots, n_b\}$ such that $\delta_k(i) = <$
- or $\delta_1, \dots, \delta_m$ are (n_a, n_b) -equal for $\{j_1, \dots, j_m\}$ and either $m \leq n_b - n_a$ or there exists $k \in \{1, \dots, m\}$ such that $\delta_k(j_k) = <$.

Lemma 2. Let $n_a \leq n_b \in \mathbb{N}$. Let $\delta_1, \dots, \delta_m$ be ordering functions. Let $\tau_{n_a}, \dots, \tau_{n_b}$ and τ'_1, \dots, τ'_m steps. Assume that for all $i \in \{n_a, \dots, n_b\}$, for all $j \in \{1, \dots, m\}$, $\delta_j(i)$ defined implies $\tau'_j \delta_j(i) \tau_i$.

We have: If there exists $\{j_1, \dots, j_m\} \subseteq \{n_a, \dots, n_b\}$ such that $\delta_1, \dots, \delta_m$ are (n_a, n_b) -equal (resp (n_a, n_b) -strict) for $\{j_1, \dots, j_m\}$ then $\{\tau'_1, \dots, \tau'_m\} \leq_m \{\tau_1, \dots, \tau_n\}$ (resp. $<_m$).

Proof. The proof is a reorganisation of the proof of [37, Lemma 14]. However, since we extended the notion and is one of the fundamental result to apply inductive lemmas, we show the proof again here.

Consider first the case where $\delta_1, \dots, \delta_m$ are (n_a, n_b) -equal for $\{j_1, \dots, j_m\}$. Thanks to Item 2, we know that the indices j_1, \dots, j_m are distinct and by Item 1 that $\delta_k(j_k)$ is defined for all $k \in \{1, \dots, m\}$. By being defined, we deduce that for all $k \in \{1, \dots, m\}$, $\tau'_k \delta_k(j_k) \tau_{j_k}$ meaning that $\tau'_k \leq \tau_{j_k}$. Hence, we have shown that the steps τ'_1, \dots, τ'_m are smaller than m distinct steps of the multiset $\{\tau_{n_a}, \dots, \tau_{n_b}\}$, which entails $\{\tau'_1, \dots, \tau'_m\} \leq_m \{\tau_{n_a}, \dots, \tau_{n_b}\}$.

When $\delta_1, \dots, \delta_m$ are (n_a, n_b) -strict for $\{j_1, \dots, j_m\}$, we do a case analysis on which item of Definition 10 is satisfied by $\delta_1, \dots, \delta_m$.

- Item 1: In such a case, we know that for all $j \in \{1, \dots, m\}$, there exists $i \in \{n_a, \dots, n_b\}$ such that $\delta_j(i) = <$ implying that $\delta_j(i)$ is defined. Thus by hypothesis, $\tau'_j < \tau_i \leq \max(\tau_{n_a}, \dots, \tau_{n_b})$. Since for all $j \in \{1, \dots, m\}$, $\tau'_j < \max(\tau_{n_a}, \dots, \tau_{n_b})$, we directly obtain that $\{\tau'_1, \dots, \tau'_m\} <_m \{\tau_{n_a}, \dots, \tau_{n_b}\}$.
- Item 2: In such a case, we know that $\delta_1, \dots, \delta_m$ are (n_a, n_b) -equal for $\{j_1, \dots, j_m\}$. Hence we already proved the steps τ_1, \dots, τ_m are smaller than m distinct steps of the multiset $\{\tau_{n_a}, \dots, \tau_{n_b}\}$. When $m \leq n_b - n_a$, it directly entails that $\{\tau'_1, \dots, \tau'_m\} <_m \{\tau_{n_a}, \dots, \tau_{n_b}\}$. When $m = n_b - n_a + 1$, the fact that there exists $k \in \{1, \dots, m\}$ such that $\delta_k(j_k) = <$ guarantees that at least one of the ordering functions is strict, and so $\tau'_k < \tau_{j_k}$. This allows us to conclude that $\{\tau'_1, \dots, \tau'_m\} <_m \{\tau_{n_a}, \dots, \tau_{n_b}\}$. \square

As previously mentioned, given a query or lemma of the form $\varrho = (F_1 \wedge \dots \wedge F_n \Rightarrow \psi)$, we will always assume the user-defined predicates occur in the premise after the standard predicates event, att and mess. Hence for each query, we can associate an integer indicating the first index where a user-defined predicate occurs in the premise. By denoting $\text{idx}(\varrho)$, we have that for all $1 \leq i < \text{idx}(\varrho)$, $\text{pred}(F_i) \in \{\text{event}, \text{att}, \text{mess}\}$, and for all $\text{idx}(\varrho) \leq i \leq n$, $\text{pred}(F_i) \in \mathcal{F}_p$. When the premise does not contain user-defined predicates, we define $\text{idx}(\varrho) = n + 1$.

Similarly, as an ordered clauses $R = (H \rightarrow \bigwedge_{i=1}^n F_i)$ originates from a query, we can similarly define the first index where a user-defined predicate occurs in $\bigwedge_{i=1}^n F_i$, denoted $\text{idx}(R)$.

Definition 11. Let $n_{\text{idx}}, n, m_{\text{idx}}, m \in \mathbb{N}$ such that $n_{\text{idx}} \leq n + 1$ and $m_{\text{idx}} \leq m + 1$. We say that the ordering functions $\delta_1, \dots, \delta_m$ are strict for (n_{idx}, n) and m_{idx} when there exist $\{j_1, \dots, j_{m_{\text{idx}}}\} \subseteq \{1, \dots, n_{\text{idx}} - 1\}$ and $\{j_{m_{\text{idx}}}, \dots, j_m\} \subseteq \{n_{\text{idx}}, \dots, n\}$ such that:

- either $\delta_1, \dots, \delta_{m_{\text{idx}}-1}$ are $(1, n_{\text{idx}} - 1)$ -strict for $\{j_1, \dots, j_{m_{\text{idx}}-1}\}$
- or $\delta_1, \dots, \delta_{m_{\text{idx}}-1}$ are $(1, n_{\text{idx}} - 1)$ -equal for $\{j_1, \dots, j_{m_{\text{idx}}-1}\}$ and $\delta_{m_{\text{idx}}}, \dots, \delta_m$ are (n_{idx}, n) -strict for $\{j_{m_{\text{idx}}}, \dots, j_m\}$

Lemma 3. Let $n_{\text{idx}}, n, m_{\text{idx}}, m \in \mathbb{N}$ such that $n_{\text{idx}} \leq n + 1$ and $m_{\text{idx}} \leq m + 1$. Let $\delta_1, \dots, \delta_m$ be strict for (n_{idx}, n) and m_{idx} . Let T be a trace.

For all tuples of steps $\tilde{\tau} = (\tau_1, \dots, \tau_{n_{\text{idx}}-1})$, $\tilde{\tau}_p = (\tau_{n_{\text{idx}}}, \dots, \tau_n)$, $\tilde{\tau}' = (\tau'_1, \dots, \tau'_{m_{\text{idx}}-1})$ and $\tilde{\tau}'_p = (\tau'_{m_{\text{idx}}}, \dots, \tau'_m)$, if for all $j \in \{1, \dots, m\}$, for all $i \in \{1, \dots, n\}$, $\delta_j(i)$ defined implies $\tau'_j \delta_j(i) \tau_i$ then $(T, \tilde{\tau}', \tilde{\tau}'_p) <_{\text{ind}} (T, \tilde{\tau}, \tilde{\tau}_p)$

Proof. For $(T, \tilde{\tau}', \tilde{\tau}'_p) <_{\text{ind}} (T, \tilde{\tau}, \tilde{\tau}_p)$ to hold, we need to prove that either $\{\tau'_1, \dots, \tau'_{m_{\text{idx}}-1}\} <_m \{\tau_1, \dots, \tau_{n_{\text{idx}}-1}\}$ or else $\{\tau'_1, \dots, \tau'_{m_{\text{idx}}-1}\} = \{\tau_1, \dots, \tau_{n_{\text{idx}}-1}\}$ and $\{\tau'_{m_{\text{idx}}}, \dots, \tau'_m\} <_m \{\tau_{n_{\text{idx}}}, \dots, \tau_n\}$.

We do a case analysis on the item satisfied in Definition 11. In Case 1, we know that $\delta_1, \dots, \delta_{m_{\text{idx}}-1}$ are $(1, n_{\text{idx}} - 1)$ -strict for $\{j_1, \dots, j_{m_{\text{idx}}-1}\}$, hence we conclude by applying Lemma 2 to obtain $\{\tau'_1, \dots, \tau'_{m_{\text{idx}}-1}\} <_m \{\tau_1, \dots, \tau_{n_{\text{idx}}-1}\}$

In Case 2, we know that $\delta_1, \dots, \delta_{m_{\text{idx}}-1}$ are $(1, n_{\text{idx}} - 1)$ -equal for $\{j_1, \dots, j_{m_{\text{idx}}-1}\}$. Hence by Lemma 2, we deduce that $\{\tau_1, \dots, \tau_{n_{\text{idx}}-1}\} \leq_m \{\tau'_1, \dots, \tau'_{m_{\text{idx}}-1}\}$. If $\{\tau_1, \dots, \tau_{n_{\text{idx}}-1}\} <_m \{\tau'_1, \dots, \tau'_{m_{\text{idx}}-1}\}$ then we conclude; otherwise $\{\tau_1, \dots, \tau_{n_{\text{idx}}-1}\} = \{\tau'_1, \dots, \tau'_{m_{\text{idx}}-1}\}$. However, since we also know that $\delta_{m_{\text{idx}}}, \dots, \delta_m$ are (n_{idx}, n) -strict for $\{j_{m_{\text{idx}}}, \dots, j_m\}$, then we can apply Lemma 2 to obtain $\{\tau'_{m_{\text{idx}}}, \dots, \tau'_m\} <_m \{\tau_{n_{\text{idx}}}, \dots, \tau_n\}$, which allows us to conclude. \square

We can now state our extended transformation rule for applying inductive lemmas, denoted $(\text{Ind}_o(\mathcal{L}_i, \mathcal{S}_p))$.

$$\frac{\begin{array}{l} \mathbb{C} \cup \{R = (H \rightarrow \bigwedge_{i=1}^n F_i)\} \\ \varrho = (\bigwedge_{i=1}^m F_i \Rightarrow \bigvee_{j=1}^k \psi_j) \in \mathcal{L}_i \\ \varrho \text{ matches } R \text{ with } \delta_1, \dots, \delta_m \text{ and } \sigma \\ \delta_1, \dots, \delta_m \text{ are strict for } (\text{idx}(R), n) \text{ and } \text{idx}(\varrho) \\ \text{for all } i, \text{pred}(F_i) \notin \mathcal{F}_p \text{ implies } \delta \sqsupseteq \delta_i \end{array}}{\mathbb{C} \cup \{H \wedge [\psi_k \sigma]_{\delta}^b \rightarrow \bigwedge_{i=1}^n F_i \bigwedge_{j=1}^k \psi_j\}}$$

Soundness of the solving procedure. The soundness result of the solving procedure is given by the following theorem that is an adaptation of [37, Theorem 4].

Theorem 4. Let \mathcal{C}_I be an initial instrumented configuration. Let \mathcal{S}_p be a set of predicates containing the predicate event, all predicates in \mathcal{F}_{bp} and all blocking predicates. Let $\mathcal{L}, \mathcal{L}_i$ be two sets of lemmas. Let $T \in \text{trace}_{IO}(\mathcal{C}_I)$. Let \mathbb{C} be a set of simplified and selection free clauses Horn clauses consistent with \mathcal{F}_p containing the selection free clauses of \mathbb{C}_{std} .

For all ordered clauses R satisfying \mathcal{S}_p , for all ordered derivations \mathcal{D} of $\bigwedge_{i=1}^m F_i$ at steps (τ_1, \dots, τ_m) from $\{R\}$ and $\mathbb{C} \cup \mathbb{C}_b(T)$ such that $\text{Hyp}_{\mathcal{L}, \mathcal{L}_i}(T, (\tau_1, \dots, \tau_{\text{idx}(R)-1}), (\tau_{\text{idx}(R)}, \dots, \tau_m))$ and $T, \mathcal{S}_p \vdash \mathcal{D}$, there exists an ordered derivation \mathcal{D}' of $\bigwedge_{i=1}^m F_i$ at steps $\tilde{\tau}$ from $\text{satrate}_{\mathcal{L}, \mathcal{L}_i}^{\mathcal{S}_p}(\{R\}, \mathbb{C})$ and $\mathbb{C} \cup \mathbb{C}_b(T)$ such that $T, \mathcal{S}_p \vdash \mathcal{D}'$.

Sketch proof. The proof can be directly adapted from the one of [37, Theorem 4]. It also relies on showing how derivations can be transformed and use produce Horn clauses generated by the transformation rules. The main change comes the proof of soundness of the application of inductive lemmas which now relies on Lemma 3 instead of [37, Lemma 14]. \square

B.5. Verification of the correspondence query

In [37, Section 7], a general algorithm is described to verify correspondence queries, including queries containing nested and injective queries. For the main query to solve, we keep these algorithms exactly as they are. The key difference comes from the way we prove attacker facts in a lemma. In a simple query, say $F \Rightarrow \text{att}(M)$, the semantics of ProVerif allows M to be deducible from the set of attacker clauses, namely $\{(RInit), (RGen), (RFail), (Rf)\}$ (see [37, Definition 30]). Thus, M may not appear in the attacker knowledge in the trace executing F . To apply our lemmas with attacker facts $\text{att}(M)$ in their conclusion, we need M to appear in the attacker knowledge before the step where F is satisfied. Otherwise, we would break the soundness of lemma simplification rule.

Thus, when proving a Lemma, instead of requesting that $\text{att}(M)$ is deducible from the restricted set of attacker clauses, we will request that $\text{att}(M)$ occurs directly in the hypothesis of the ordered clauses. Thanks to Item 2 of Definition 6, this will ensure that M is already in the attacker knowledge.

As an interesting consequence, ProVerif may fail to prove a correspondence query when it is declared as a lemma but succeed when it is declared as a standard query.