Automatically Checking Commitment Protocols in ProVerif without False Attacks

Tom Chothia[1], Ben Smyth[2], and Chris Staite[1]

[1] School of Computer Science, University of Birmingham, UK
[2] Mathematical and Algorithmic Sciences Lab, France Research Center,
Huawei Technologies Co. Ltd., France

Abstract. ProVerif over-approximates the attacker's power to enable verification of processes under replication. Unfortunately, this results in ProVerif finding false attacks. This problem is particularly common in protocols whereby a participant commits to a particular value and later reveals their value. We introduce a method to reduce false attacks when analysing secrecy. First, we show how inserting phases into non-replicated processes enables a more accurate translation to Horn clauses which avoids some false attacks. Secondly, we generalise our methodology to processes under replication. Finally, we demonstrate the applicability of our technique by analysing BlueTooth Simple Pairing. Moreover, we propose a simplification of this protocol that achieves the same security goal.

1 Introduction

State space exploration has emerged as a leading verification technique [25] and, in this context, Abadi & Fournet [2] propose the applied pi calculus – an extension of the pi calculus – to reason with cryptographic protocols. Unfortunately, proving security in this context is undecidable [21], due to several sources of unboundedness, including, messages of arbitrary length and the possibility of an unbounded number of sessions. Accordingly, state-of-the-art automated reasoning techniques focus on sound, but incomplete, methodologies, which may report false attacks and do not always terminate.

Blanchet [8–10] translates applied pi calculus processes to Horn clauses and uses resolution of Horn clauses to reason with secrecy and authentication properties, these results have been implemented in ProVerif [15]. ProVerif has been successfully used to automatically analyse cryptographic protocols from a variety of applications domains, including, key exchange [1, 10, 29], electronic voting [20, 5, 4] and trusted computing [18, 19, 28], for example. However, Blanchet's translation to Horn clauses over-approximates the attacker's power and, therefore, ProVerif may report false attacks; as highlighted by Blanchet [12, §2.2]:

---

[*] A long version of this paper and ProVerif source code supporting this paper are available from `http://www.cs.bham.ac.uk/~tpc/projects/falseattacks`.

> *"false attacks occur typically for protocols that first need to keep data secret, then publish them later in the protocol. In that situation, the Horn clause model considers that the attacker can re-inject the secret in the early part of the run, which is not possible in reality."*

This behaviour is typical of protocols in which participants commit to a value and later reveal it, such as the BlueTooth Simple Pairing protocol [22, 23], which we analyse in Section 6. In this paper, we introduce techniques for analysing secrecy that avoids some false attacks.

1.1   Our contribution in context of existing work

Let us consider the following process, proposed by Allamigeon & Blanchet [3, §3.2]:

$$\nu\, n.c(x).\bar{c}\langle n\rangle.\text{if } x = n \text{ then } \bar{c}\langle s\rangle \tag{P1}$$

Process P1 generates a fresh bound name n, binds a message input to variable x, outputs the name n, tests if the message bound to x is equal to n, and outputs the free name s, if the test succeeds. It follows intuitively that an attacker that does not know s in advance, cannot derive s from Process P1, that is, we have secrecy$(\{c\}, s)$: P1. However, the Horn clauses generated by Blanchet's translation of P1 include:

$$\texttt{attacker}(x) \Rightarrow \texttt{attacker}(n)$$
$$\texttt{attacker}(n) \Rightarrow \texttt{attacker}(s)$$

Hence, ProVerif cannot prove secrecy$(\{c\}, s)$: P1, because the Horn clauses permit the following false attack: knowledge of an arbitrary term M implies knowledge of n and knowledge of n implies knowledge of s. This is due to an over-approximation: the Horn clauses model the process $\nu\, n.!c(x).\bar{c}\langle n\rangle.\text{if } x = n \text{ then } \bar{c}\langle s\rangle$, rather than $P1$. It follows that the Horn clauses do not enforce that a message input must be received *before* the bound name n is output. Intuitively, such false attacks can be avoided by ensuring that the translation preserves temporal order of message inputs and outputs. We shall achieve this objective using Blanchet, Abadi & Fournet's notion of *phases* [13, §8]. Phases, denoted $t\!:\!P$, ensure which parts of concurrent processes are active at a particular time.

Inserting phases into a process can stop false attack, e.g., we can add a phase into process P1:

$$0\!:\!\nu\, n.c(x).1\!:\!\bar{c}\langle n\rangle.\text{if } x = n \text{ then } \bar{c}\langle s\rangle \tag{P2}$$

The semantics of phases ensure that $P2$ is a sound approximation of $P1$. However, the Horn clauses generated by Blanchet's translation of P2 to Horn clauses are more precise:

$$\texttt{attacker}(x) \Rightarrow \texttt{attacker}'(x)$$
$$\texttt{attacker}(x) \Rightarrow \texttt{attacker}'(n)$$
$$\texttt{attacker}(n) \Rightarrow \texttt{attacker}'(s)$$

Indeed, ProVerif can use these Horn clauses to prove $secrecy(\{c\}, s) : \text{P2}$ and, since $P2$ is a sound approximation of $P1$, we have $secrecy(\{c\}, s) : \text{P1}$.

We define a compiler (Section 3) that inserts phases into a restricted class of processes such that Blanchet's translation from compiled processes to Horn clauses enforces order. We prove the soundness of our methodology:

$$\text{secrecy holds in the original process}$$
$$\textit{iff} \text{ secrecy holds in the compiled process}$$

This technique is limited to proving secrecy of names which are not under the scope of a replication and we overcome this limitation using an alternative notion of secrecy.

The class of secrecy properties which can be considered using Blanchet's definition [10, Section 3.1] is limited, for example, we cannot consider secrecy of an arbitrary session secret s in the process $!\nu s.P$, that is, we cannot consider if an instance of $\nu s.P$ leaks s. This problem can be overcome by abstraction, in particular, Blanchet [10, §2.3] proposes the following solution. Extend the set of function symbols with the binary constructor e and the binary destructor d, let the set of rewrite rules $def(d) = \{d(x, e(x, y)) \rightarrow y\}$, and modify $!\nu s.P$ such that $e(s, m)$ is published at the end of every *successful* session of $\nu s.P$, where m is a free name not known by the attacker. It follows that the modified process preserves secrecy of m iff $!\nu s.P$ preserves secrecy of s, hence, we have a methodology to consider the secrecy of bound names.

We consider (Section 4) a definition for secrecy of bound names which does not require abstraction and introduce a new proof technique: given a process P, name s, and fresh name s', we have

secrecy of the bound name s in $!\nu s.P$
  $\textit{iff}$ secrecy of the bound name s' in $\nu s'.(P\{s'/s\}) \mid !\nu s.P)$
  $\textit{iff}$ secrecy of the free name s' in $P\{s'/s\} \mid !\nu s.P$

In the context of these results, we describe how our compiler can be applied to avoid false attacks, in particular, we can prove security results for secrets under the scope of a replication.

We demonstrate the applicability of our technique by analysing three protocols (Sections 5 & 6): a toy extension of the Needham-Schroeder protocol in which one of the participants reveals their nonce at the end of a successful run, the Bluetooth Simple Pairing [22, 23], and a simplification of the Bluetooth Simple Pairing that we propose. ProVerif finds false attacks against each of these protocols, whereas our techniques allow us to prove security.

Pairing protocols typically use a low-entropy, human-verifiable string, derived from a high-entropy shared secret, to authenticate protocol participants and protect against impersonation attacks. Given that the string is low-entropy, strings derived from distinct secrets may collide and an attacker that can predict collisions can launch impersonation attacks. Accordingly, pairing protocols must

ensure that deriving collisions is computationally expensive. The Bluetooth Simple Pairing protocol defends against such attacks by making both parties commit to particular values before the low-entropy string can be calculated; this deprives an attacker of the opportunity to carry out a brute force attack to find a collision.

We develop theory to enable protocols that are vulnerable to collision attacks to be accurately modelled in the applied pi calculus. Accordingly, our analysis of Bluetooth Simple Pairing is more precise than an earlier analysis by Chang & Shmatikov [17], which ignores collision attacks, and so would incorrectly find the protocol to be secure even if the steps that stop impersonation attacks were removed. We also present a simplified version of the Bluetooth Simple Pairing protocol, which achieves the same secrecy goals with fewer steps, and we use our analysis method to show that it is secure.

The key contributions of this paper are:

- A framework for avoiding some false attacks when analysing secrecy.
- A definition for secrecy of bound names.
- A method to captured collision attacks.
- A demonstration of how BlueTooth Simple Pairing defends against collision attacks.
- A simplified pairing protocol.

Hence, our paper advances automated analysis techniques.

2   Background: Applied pi calculus

We adopt Blanchet's dialect [10] of the applied pi calculus [2, 27], which is suited to automated reasoning using Blanchet's ProVerif [15]. The dialect uses the notion of configurations proposed by Baudet [7] to avoid structural equivalence, which simplifies security definitions and subsequent proofs.

The calculus assumes an infinite set of *names*, an infinite set of *variables*, and a finite set of *function symbols* (*constructors* and *destructors*), each with an associated arity. We write f for a constructor, g for a destructor; constructors are used to build terms whereas destructors are used to manipulate terms in processes. *Terms* range over names, variables, and applications of constructors to terms. *Substitutions* $\{M/x\}$ replace the variable x with the term M. Arbitrarily large substitutions can be written as $\{M_1/x_1, \ldots, M_n/x_n\}$ and the letters σ and τ range over substitutions. We write $M\sigma$ for the result of applying σ to the variables of M.

The signature Σ is equipped with a finite set of equations of the form $M = N$ and we derive an equational theory from this set by reflexive, symmetric and transitive closure, closure under the application of constructors, closure under substitution of terms for variables, and closure under bijective renaming of names. We write $\Sigma \vdash M = N$ for an equality modulo the equational theory and $\Sigma \vdash M \neq N$ for an inequality modulo the equational theory. (We write $M = N$ and $M \neq N$ for syntactic equality and inequality, respectively.)

The semantics of a destructor g of arity l is given by a finite set $\text{def}(g)$ of rewrite rules $g(M'_1, \ldots, M'_l) \to M'$, where $M'_1, \ldots, M'_l, M'$ are terms that contain only constructors and variables; the variables of M' must be bound in $M'_1, \ldots, M'_l$ and variables are subject to renaming. The value $g(M_1, \ldots, M_l)$ is defined if and only if there exists a substitution σ and a rewrite rule $g(M'_1, \ldots, M'_l) \to M'$ in $\text{def}(g)$ such that $M_i = M'_i \sigma$ for all $i \in \{1, \ldots, l\}$, and in this case $g(M_1, \ldots, M_l)$ is defined as $M'\sigma$.

The grammar for terms and *processes* is presented in Figure 1, where t is a non-negative integer representing a global clock. The process *let* $x = g(M_1, ..., M_l)$ *in P else Q* tries to evaluate $g(M_1, ..., M_l)$; if this succeeds (that is, if $g(M_1, ..., M_l)$ is defined), then x is bound to the result and P is executed, otherwise, Q is executed. The statement let $x = g(M_1, \ldots, M_l)$ in P else Q may be abbreviated as let $x = g(M_1, \ldots, M_l)$ in P, when Q is 0. The syntax does not include the conditional if $M = N$ then P else Q, but this can be defined as let $x = \mathsf{eq}(M, N)$ in P else Q, where x is a fresh variable, eq is a binary destructor, and $\text{def}(\mathsf{eq}) = \{\mathsf{eq}(x, x) \to x\}$; we always include eq in our set of function symbols. For convenience, we may write if $M = N$ then P else Q for let $x = \mathsf{eq}(M, N)$ in P else Q and if $M = N$ then P for let $x = \mathsf{eq}(M, N)$ in P. In Figure 1, we extend Blanchet's syntax [10] with Blanchet, Abadi & Fournet's notion of *phases* [13, §8], denoted $t\!:\!P$, which ensures a process $t\!:\!P$ is only active during time t.

Fig. 1 Syntax for terms and processes

$M, N ::=$	terms
x, y, z	variables
a, b, c, k, s	names
$f(M_1, \ldots, M_n)$	constructor application
$D ::= g(M_1, \ldots, M_n)$	destructor application
$P, Q ::=$	processes
0	nil
$\overline{M}\langle N \rangle.P$	output
$M(x).P$	input
$P \mid Q$	parallel composition
$!P$	replication
$\nu\, a.P$	restriction
let $x = D$ in P else Q	term evaluation
$t\!:\!P$	phase

The sets of free and bound names, respectively variables, in process P are denoted by $\text{fn}(P)$ and $\text{bn}(P)$, respectively $\text{fv}(P)$ and $\text{bv}(P)$. We also write $\text{fn}(M)$ and $\text{fv}(M)$ for the sets of names and variables in term M. A process P is closed if it has no free variables. A context C is a process with a hole and we obtain $C[P]$

as the result of filling C's hole with P. An evaluation context is a context whose hole is not in the scope of a replication, an input, an output, or a term evaluation.

The operational semantics (Figures 2) for the applied pi-calculus are defined by reduction ($\rightarrow$) on *configurations*. A configuration $\mathcal{C}$ is a pair $E, \mathcal{P}$ such that E is a finite set of names, and $\mathcal{P}$ is a finite multiset of pairs of closed process. The set E contains all the free names in $\mathcal{P}$, and is extended to include any names introduced during reduction, namely, those names introduced by $(E, \mathcal{P} \cup \{t : \nu a; P\}) \rightarrow (E \cup \{a'\}, \mathcal{P} \cup \{t : P\{a'/a\}\})$. A sequence of reductions, denoted $\mathcal{C}_1 \rightarrow \mathcal{C}_2 \rightarrow \cdots \rightarrow \mathcal{C}_n$, is called a *trace*. We occassionally write $\rightarrow^*$ for the reflexive and transitive closure of $\rightarrow$.

Fig. 2 Operational semantics

$E, \mathcal{P} \cup \{t{:}0\} \rightarrow E, \mathcal{P}$ $\hfill$ (RED NIL)

$E, \mathcal{P} \cup \{t{:}!P\} \rightarrow E, \mathcal{P} \cup \{t{:}!P, t{:}P\}$ $\hfill$ (RED REPL)

$E, \mathcal{P} \cup \{t{:}(P \mid Q)\}) \rightarrow E, \mathcal{P} \cup \{t{:}P, t{:}Q\}$ $\hfill$ (RED PAR)

$E, \mathcal{P} \cup \{t{:}\nu a.P\} \rightarrow E \cup \{a'\}, \mathcal{P} \cup \{t{:}P\{a'/a\}\}$ $\hfill$ (RED RES)
for some name $a' \notin E$

$E, \mathcal{P} \cup \{t{:}\overline{N}\langle M \rangle.P, t{:}N(x).Q\} \rightarrow E, \mathcal{P} \cup \{t{:}P, t{:}Q\{M/x\}\}$ $\hfill$ (RED I/O)

$E, \mathcal{P} \cup \{t{:}\text{let } x = D \text{ in } L\} \rightarrow E, \mathcal{P} \cup \{t{:}P\{M/x\}\}$ $\hfill$ (RED DESTR 1)
if there exists M such that $D \rightarrow M$

$E, \mathcal{P} \cup \{t{:}\text{let } x = D \text{ in } P \text{ else } Q\}) \rightarrow E, \mathcal{P} \cup \{t{:}Q\}$ $\hfill$ (RED DESTR 2)
if there is no M such that $D \rightarrow M'$

$E, \mathcal{P} \cup \{t{:}t'{:}P\} \rightarrow E, \mathcal{P} \cup \{t'{:}P\}$ $\hfill$ (RED ORDER)
if $t < t'$

Given a process P in the applied pi-calculus without phases and a set of names *Init*, the configuration *Init*, $\{0{:}P\}$ will reduce using our semantics in exactly the same way as the configuration *Init*, $\{P\}$ using Blanchet's semantics [10]. We note that Blanchet, Abadi & Fournet [13, §8], who introduce phases, assume any process without phases is assumed to run in phase zero. In this paper, we wish to distinguish between processes in phase zero and processes without phases, so we always make the phases explicit, except in case studies where we adopt Blanchet, Abadi & Fournet's convention for brevity.

3   Secrecy of free names

We recall (Definition 1) Blanchet's formalisation [10] of knowledge derivable from a trace – that is, a reduction on a configuration – as any names which are output.

Definition 1. *Let $\mathcal{T} = E_0, \mathcal{P}_0 \rightarrow^* E', \mathcal{P}'$ be a trace, n be a name, and Init be a finite set of names. We write* attacker$(Init, n) : \mathcal{T}$ *if $\mathcal{T}$ contains a reduction $E, \mathcal{P} \cup \{t : \bar{c}\langle n \rangle.P, t : c(x).Q\} \rightarrow E, \mathcal{P} \cup \{t : P, t : Q\{n/x\}\}$ for some $E, \mathcal{P}, P, Q, x, t$ and $c \in Init$.*

It follows naturally that a configuration preserves the secrecy of a name if no (adversarial) process added to the configuration can generate a trace which permits the name to be derived. We recall (Definition 2) Blanchet's definition [10] for secrecy of free names.

Definition 2. *Let $\mathcal{C} = E, \mathcal{P}$ be a configuration, n be a name and Init be a finite set of names, where $\mathrm{fn}(\mathcal{P}) \subseteq E$. We write* secrecy$(Init, n) : \mathcal{C}$ *if for all processes Q such that $\mathrm{fn}(Q) \subseteq Init$ there is no trace $\mathcal{T} = E \cup Init \cup \mathrm{fn}(n), \mathcal{P} \cup \{Q\} \rightarrow^* E', \mathcal{P}'$ such that* attacker$(Init, n) : \mathcal{T}$ *for some E' and $\mathcal{P}'$.*

Let P be a closed process, n be a name, and Init be a finite set of names. We write secrecy$(Init, n) : P$ *if* secrecy$(Init, n) : \mathrm{fn}(P), \{P\}$.

Definition 2 facilitates the analysis of secrecy when the secret is a free name and Section 4 proposes a definition which supposes the secret is bound.

As discussed in Section 1.1, we encounter false attacks when analysing secrecy of free names and we overcome this problem in the remainder of this section.

3.1   Our compiler: Phases improve Horn clause generation

As demonstrated in Section 1.1, false attacks can be avoided by inserting phases into processes. Formally, we insert phases using function δ:

Definition 3. *Given a set of names Init and a process, we define δ as follows:*

$$\delta(Init, P) = \{P\} \cup \{C[1 : M(x).P'] : P = C[M(x).P']\}$$
$$\cup \{C[1 : \overline{M}\langle N \rangle.P'] : P = C[\overline{M}\langle N \rangle.P'] \wedge \mathrm{fn}(M) \cap Init = \emptyset\}$$

Function δ outputs a set of processes representing all ways of inserting a phase into P such that the phase appears immediately before an input or an output on a private channel[1].

The insertion of one phase does not generally result in a sound abstraction. For instance, process $P2 \mid 0 : !P1$ is not a sound approximation of $0 : P1 \mid 0 : !P1$, because the phase in $P2$ prevents the instance of n generated by $P2$ being input by $0 : P1$. This problem can be overcome by ensuring that all inputs and outputs are available in either phase 0 or phase 1, which can be achieved using our compiler (Definition 4). For simplicity, we restrict our compiler to multisets of processes $\mathcal{P} = \{0 : L, 0 : !L_1, \ldots, 0 : !L_m\}$, where $L, L_1, \ldots, L_m$ are *linear processes* (Figure 3) and ProVerif discovers false attacks arising from process L. This is sufficient to avoid the false attacks discovered in the examples in this paper.

---

[1] Outputs on public channels can be received by the environment in phase 0 and replayed in phase 1, therefore function δ does not insert phases immediately before outputs on public channels.

Fig. 3 Syntax for linear processes

$L ::=$	linear processes
$\quad 0$	nil
$\quad \overline{M}\langle N\rangle.L$	output
$\quad M(x).L$	input
$\quad \nu\, a.L$	restriction
$\quad$ let $x = g(M_1,\ldots,M_n)$ in L else 0	destructor application

---

Definition 4. *Given a set of names Init and a multiset of process $\mathcal{P} = \{0 : L, 0\!:\!!L_1, \ldots, 0\!:\!!L_m\}$, we define Δ as follows:*

$$\Delta(Init, \mathcal{P}) = \left\{ \{P\} \cup \mathcal{Q} \ \middle| \ P \in \delta(Init, 0\!:\!L) \wedge \mathcal{Q} = \bigcup_{1 \leq i \leq m} \delta(Init, 0\!:\!!L_i) \right\}$$

The compiler tries to avoid false attacks by inserting a phase into process L (in a different place for each member of the set produced by Δ). To ensure that all of the original reductions are still possible, our compiler also generates a copy of every other process with a phase in every necessary position.

3.2   Automated reasoning without false attacks

Our compiler is designed such that Blanchet's translation from compiled processes to Horn clauses ensures that the clauses abide by an ordering, thereby avoiding the false attacks described in Section 1.1, whilst preserving secrecy:

Theorem 1. *Given a name s, sets of names E and Init and a multiset of processes $\mathcal{P} = \{0\!:\!L, 0\!:\!!L_1, \ldots, 0\!:\!!L_m\}$, such that $s \notin \bigcup_{1 \leq i \leq n}(\mathrm{fn}(L_i) \cup \mathrm{bn}(L_i))$ and $\mathrm{fn}(\mathcal{P}) \subseteq E$, we have for all $\mathcal{Q} \in \Delta(Init, \mathcal{P})$ that:*

$$\mathrm{secrecy}(Init, s) : E, \mathcal{P} \Leftrightarrow \mathrm{secrecy}(Init, s) : E, \mathcal{Q}$$

The proof of Theorem 1 appears in the long version of this paper. We demonstrate an application of Theorem 1 with reference to processes $P1$ and $P2$:

Example 1. Let $C = 0 : \nu\, n.c(x).\_$ and witness that $0 : P1 = C[\overline{c}\langle n\rangle.\text{if } x = n \text{ then } \overline{c}\langle s\rangle]$ and $P2 = C[1 : \overline{c}\langle n\rangle.\text{if } x = n \text{ then } \overline{c}\langle s\rangle]$, i.e., $\{P2\} \in \Delta(\{c\}, \{P1\})$. We have $\mathrm{secrecy}(\{c\}, s) : \{c, s\}, \{0\!:\!P1\} \Leftrightarrow \mathrm{secrecy}(\{c\}, s) : \{c, s\}, \{P2\}$ by Theorem 1, hence, $\mathrm{secrecy}(\{c\}, s) : (0\!:\!P1) \Leftrightarrow \mathrm{secrecy}(\{c\}, s) : P2$ by Definition 2. Moreover, since ProVerif can prove $\mathrm{secrecy}(\{c\}, s) : P2$, we have the desired result, namely $\mathrm{secrecy}(\{c\}, s) : (0\!:\!P1)$.

4   Secrecy of bound names

We adapt the notion of knowledge derivable from a trace (Definition 1) to consider names which are bound by the trace's initial configuration.

Definition 5. *Let $\mathcal{T} = E_0, \mathcal{P}_0 \rightarrow^* E_1, \mathcal{P}_1$ be a trace, n be a name, and Init be a finite set of names. We write $\widehat{\text{attacker}}(Init, n) : \mathcal{T}$ if $\mathcal{T}$ contains the following reductions*

$$
\begin{aligned}
E, \mathcal{P} \cup \{t{:}\nu\, n.P\} &\rightarrow E \cup \{n'\}, \mathcal{P} \cup \{t{:}P\{n'/n\}\} \\
&\rightarrow^* E', \mathcal{P}' \cup \{s{:}\bar{c}\langle n'\rangle.P', s{:}c(x).Q'\} \\
&\rightarrow E', \mathcal{P}' \cup \{s{:}P', s{:}Q'\{n'/x\}\}
\end{aligned}
$$

where $c \in Init$ and some $E, E', \mathcal{P}, \mathcal{P}', P, P', Q', n', x, s$ and t.

Intuitively, a trace $\mathcal{T}$ satisfies $\widehat{\text{attacker}}(Init, n) : \mathcal{T}$ if n is bound by the trace's initial configuration, n is renamed to n' and, subsequently, n' is outputted.

It follows naturally from Definition 5 that a configuration preserves secrecy of a bound name if no (adversarial) process added to the configuration can generate a trace which permits the name to be derived.

Definition 6. *Let $\mathcal{C} = E, \mathcal{P}$ be a configuration, n be a name and Init be a set of names, where $\text{fn}(\mathcal{P}) \subseteq E$. We write $\widehat{\text{secrecy}}(Init, n) : \mathcal{C}$ if for all processes Q such that $\text{fn}(Q) \subseteq Init$ and $n \notin \text{bn}(Q)$, there is no trace $\mathcal{T} = E \cup Init, \mathcal{P} \cup \{Q\} \rightarrow^* E', \mathcal{P}'$ such that $\widehat{\text{attacker}}(Init, n) : \mathcal{T}$ for some E' and $\mathcal{P}'$.*

Let P be a closed process, n be a name and Init be a finite set of names, where $n \in \text{bn}(P)$. We write $\widehat{\text{secrecy}}(Init, n) : P$ if $\widehat{\text{secrecy}}(Init, n) : \text{fn}(P), \{P\}$.

We remark that $\widehat{\text{secrecy}}(Init, n) : P$ guarantees secrecy of every bound name n in P. It follows that $\widehat{\text{secrecy}}(Init, n) : (s{:}\nu\, n.Q \mid t{:}\nu\, n.R)$ implies $\widehat{\text{secrecy}}(Init, n) : (s{:}\nu\, n.Q \mid t{:}\nu\, m.(R\{m/n\}) \wedge \widehat{\text{secrecy}}(Init, m) : (s{:}\nu\, n.Q \mid t{:}\nu\, m.(R\{m/n\})))$, for example.

Secrecy of bound names is not new. ProVerif can already check the secrecy of bound names, however, the corresponding theoretical definition has not been published. (Blanchet [11] has confirmed that Definition 6 corresponds to the secrecy of bound names notion used by ProVerif.) In addition, Ryan & Smyth [27, §3.1] propose a definition for secrecy of bound names in the applied pi calculus, however, their definition is restricted to bound names which do not appear under the scope of replication and we do not impose such a restriction. The false attacks we encounter when analysing secrecy of free names similarly occur when analysing secrecy of bound names, for instance, ProVerif cannot prove $\widehat{\text{secrecy}}(\{c\}, s) : \{c\}, \{0{:}!\nu\, s.P1\}$. In the remainder of this section we overcome this problem.

4.1 A proof technique for secrecy of bound names

It follows from our semantics and the definition of bound secrecy that: if $t{:}!\nu\, n.P$ does not preserve the secrecy of n, then there exists an instance of $\nu\, n.P$ that leaks n, when running in parallel with $!\nu\, n.P$. We now show that it is sufficient to rename n with some fresh name m in an instance of $\nu\, n.P$ and consider secrecy of the m in $t{:}\nu\, m.(P\{m/n\}) \mid !\nu\, n.P)$.

Theorem 2. *Given a name s, sets of names E and $Init$, process P and multiset of processes $\mathcal{P}$ such that $s \notin \mathrm{bn}(P) \cup \mathrm{bn}(\mathcal{P})$ and $\mathrm{fn}(\mathcal{P}) \cup (\mathrm{fn}(P) \setminus \{s\}) \subseteq E$, we have, for all fresh names s', that:*

$$\widehat{\mathrm{secrecy}}(Init, s) : E, \mathcal{P} \cup \{t : !\nu\, s.P\} \iff$$
$$\widehat{\mathrm{secrecy}}(Init, s') : E, \mathcal{P} \cup \{t : !\nu\, s.P, t : \nu\, s'.(P\{s'/s\})\}$$

The proof of Theorem 1 appears in the long version of this paper.

4.2 Secrecy of bound and free names coincide

Secrecy of bound and free names coincide when the secret is not under replication.

Proposition 1. *Given a name s, sets of names E and $Init$, process P and multiset of processes $\mathcal{P}$ such that $s \notin \mathrm{bn}(P) \cup \mathrm{bn}(\mathcal{P}) \cup \mathrm{fn}(\mathcal{P}) \cup Init \cup E$ and $\mathrm{fn}(\mathcal{P}) \cup (\mathrm{fn}(P) \setminus \{s\}) \subseteq E$, we have:*

$$\widehat{\mathrm{secrecy}}(Init, s) : E, \mathcal{P} \cup \{t : \nu\, s.P\} \iff \mathrm{secrecy}(Init, s) : E \cup \{s\}, \mathcal{P} \cup \{t : P\}$$

Proof sketch. We note that $E, \mathcal{P} \cup \{t : \nu\, s.P\} \to E \cup \{s'\}, \mathcal{P} \cup \{t : P\{s'/s\}\}$, for a fresh name s' and that $\mathrm{secrecy}(Init, s) : E \cup \{s\}, \mathcal{P} \cup \{t : P\}$ iff $\mathrm{secrecy}(Init, s') : E \cup \{s'\}, \mathcal{P} \cup \{t : P\{s'/s\}\}$. Therefore, given a trace that causes free secrecy to fail to hold for the R.H.S we can add the new name declaration to produce a trace that causes secrecy to fail for the L.H.S. Conversely, given a trace that causes secrecy to fail for the L.H.S., removing the new name declaration will give us a trace that causes secrecy to fail for the R.H.S. $\qquad\square$

A similar equivalence does not hold when the secret is under replication, as the following example demonstrates:

Example 2. For instance, suppose $P = \overline{a}\langle s\rangle \mid a(x).a(y).\mathrm{if}\ x = y\ \mathrm{then}\ \overline{c}\langle s\rangle$ and $Init = \{c\}$, we have $\widehat{\mathrm{secrecy}}(Init, s) : (0 : !\nu\, s.P)$ but not $\mathrm{secrecy}(Init, s) : (0 : !P)$, since:

$$
\begin{aligned}
E, \{0 : !P, 0 : c(z)\} &\to\to E, \{0 : !P, 0 : P, 0 : P, 0 : c(z)\} &&\text{by (RED REPL)} \\
&\to\to E, \{0 : !P, 0 : \overline{a}\langle s\rangle, 0 : \overline{a}\langle s\rangle, Q, Q, 0 : c(z)\} &&\text{by (RED PAR)} \\
&\to\to E, \mathcal{P} \cup \{0 : \mathrm{if}\ s = s\ \mathrm{then}\ \overline{c}\langle s\rangle, 0 : c(z)\} &&\text{by (RED I/O)} \\
&\to E, \mathcal{P} \cup \{0 : \overline{c}\langle s\rangle, 0 : c(z)\} &&\text{by (RED DESTR 1)} \\
&\to E, \mathcal{P} \cup \{0 : 0, 0 : 0\{s/z\}\} &&\text{by (RED I/O)}
\end{aligned}
$$

where $E = \{a, c, s\}$, $Q = 0 : a(x).a(y).\mathrm{if}\ x = y\ \mathrm{then}\ \overline{c}\langle s\rangle$ and $\mathcal{P} = \{0 : !P, 0, 0, Q\}$.

4.3 Automated reasoning without false attacks

The following corollary allows us to reduce the false attacks encountered when analysing secrecy of bound names.

Corollary 1. *Given a name s, sets of names E and Init, and a multiset of processes $\mathcal{P} = \{0\!:\!!\nu\, s.L, 0\!:\!!L_1, \ldots, 0\!:\!!L_m\}$, such that $s \notin \mathrm{bn}(L) \cup \bigcup_{1 \le i \le n}(\mathrm{fn}(L_i) \cup \mathrm{bn}(L_i))$ and $\mathrm{fn}(\mathcal{P}) \subseteq E$, we have for all fresh names s' and $\mathcal{Q} \in \overline{\Delta}(Init, \mathcal{P} \cup \{L\{s'/s\}\})$ that:*

$$\widehat{\mathrm{secrecy}}(Init, s) : E, \mathcal{P} \Leftrightarrow \mathrm{secrecy}(Init, s') : E \cup \{s'\}, \mathcal{Q}$$

Proof. We have:

$\widehat{\mathrm{secrecy}}(Init, s) : E, \mathcal{P}$
$\Leftrightarrow \widehat{\mathrm{secrecy}}(Init, s') : E, \{0\!:\!\nu\, s'.(L\{s'/s\})\} \cup \mathcal{P}$ by Theorem 2
$\Leftrightarrow \mathrm{secrecy}(Init, s') : E \cup \{s'\}, \{0\!:\!L\{s'/s\}\} \cup \mathcal{P}$ by Proposition 1 $\square$
$\Leftrightarrow \mathrm{secrecy}(Init, s') : E \cup \{s'\}, \mathcal{Q}$ by Theorem 1

We demonstrate an application of Corollary 1 by evaluating $\widehat{\mathrm{secrecy}}(\{c\}, s) : \{c\}, \{!\nu\, s.P1\}$:

Example 3. Witness that $\nu\, s.P1 = C[c(x).\bar{c}\langle n\rangle.\text{if } x = n \text{ then } \bar{c}\langle s\rangle]$, where $C = \nu\, n.\_$, but there is no other context C, process L and terms M and x such that $\nu\, s.P1 = C[M(x).L]$. In addition, there is no context C, process L and terms M, N such that $\nu\, s.P1 = C[\overline{M}\langle N\rangle.L] \wedge \mathrm{fn}(M) \cap \{c\} = \emptyset$. It follows that $\delta(\{c\}, \nu\, s.0 : P1) = \{0\!:\!!\nu\, s.P1, 0\!:\!!\nu\, s.1\!:\!c(x).\bar{c}\langle n\rangle.\text{if } x = n \text{ then } \bar{c}\langle s\rangle\}$. Let $\mathcal{Q} = \delta(\{c\}, 0\!:\!\nu\, s.P1) \cup \{P3\}$, where P3 is defined as follows:

$$0\!:\!\nu\, n.c(x).1\!:\!\bar{c}\langle n\rangle.\text{if } x = n \text{ then } \bar{c}\langle s'\rangle \tag{P3}$$

Since $P1\{s'/s\} = C[\bar{c}\langle n\rangle.\text{if } x = n \text{ then } \bar{c}\langle s'\rangle]$, where $C = \nu\, n.c(x).\_$, we have $P3 \in \delta(\emptyset, 0\!:\!P1\{s'/s\})$ and it follows that $\mathcal{Q} \in \Delta(\{c\}, \{0\!:\!P1\{s'/s\}, 0\!:\!!\nu\, s.P1\})$. We have $\widehat{\mathrm{secrecy}}(\{c\}, s) : \{c\}, \{0 : !\nu\, s.P1\} \Leftrightarrow \mathrm{secrecy}(\{c\}, s') : \{c, s'\}, \mathcal{Q}$ by Corollary 1 and, since ProVerif can prove $\mathrm{secrecy}(\{c\}, s) : \mathcal{Q}$, we have the desired result $\widehat{\mathrm{secrecy}}(\{c\}, s) : \{c\}, \{!\nu\, s.P1\}$.

5   Case study I: Needham-Schroeder protocol

For our first case study, we compose the Needham-Schroeder protocol [26] with a symmetric encryption scheme to derive a secure channel.

Needham-Schroeder protocol. The Needham-Schroeder protocol [26] is intended to generate a session key shared between two participants. We assume the participants are Alice and Bob, to prevent Lowe's man-in-the-middle attack[2] [24]. The protocol proceeds as follows. First, Alice generates a nonce and outputs

---

[2] Lowe's attack works as follows: an attacker engages Alice in a session of the protocol and impersonates Alice to Bob in a parallel session. The attack can be thwarted by assuming that Alice will only run the protocol with Bob, rather than any other principal.

the nonce encrypted with Bob's public key. Secondly, Bob decrypts Alice's ciphertext using his private key to recover Alice's nonce, generates a nonce, and outputs the pair of nonces encrypted with Alice's public key. Finally, Alice decrypts Bob's ciphertext using her private key to recover the pair of nonces and outputs Bob's nonce encrypted with Bob's public key. The rationale behind the protocol is that: since only Bob can recover Alice's nonce, only he can output the encrypted pair of nonces, moreover, since only Alice can recover Bob's nonce, only she can output his encrypted nonce; it follows that the two nonces are only known to Alice and Bob.

Symmetric encryption scheme. Symmetric encryption enables a secret to be shared between two participants. We consider the following symmetric encryption scheme. Given an identifier, key, and secret as input, the initiator outputs the identifier paired with the secret encrypted with the key. Upon receipt of such a pair, the interlocutor uses the key associated with the initiator's identifier to decrypt the ciphertext. It follows from our description that the initiator's secret can only be known by the key holders.

We compose the Needham-Schroeder protocol and symmetric encryption scheme to derive a secure channel[3]. The Needham-Schroeder protocol is used to generate a pair of nonces and these nonces are used by the symmetric encryption scheme as follows: Alice's nonce is used as the identifier and Bob's nonce is used as the key. Intuitively, the composition ensures that Alice's secret is known only to Alice and Bob, because Bob's nonce is only known to Alice and Bob.

5.1 Applied pi calculus model

We construct a signature Σ to capture the primitives modelling cryptographic operators and constants: $\Sigma = \{\mathsf{fst}, \mathsf{snd}, \mathsf{pk}, \mathsf{pair}, \mathsf{adec}, \mathsf{aenc}, \mathsf{sdec}, \mathsf{senc}\}$, where $\mathsf{fst}, \mathsf{snd}, \mathsf{pk}$ are unary functions and $\mathsf{adec}, \mathsf{aenc}, \mathsf{sdec}, \mathsf{senc}$ are binary functions. We equip the signature with the following rewrite rules:

$$\{\mathsf{fst}(\mathsf{pair}(x,y)) \to x, \ \mathsf{snd}(\mathsf{pair}(x,y)) \to y,$$
$$\mathsf{adec}(x, \mathsf{aenc}(\mathsf{pk}(x),y)) \to y, \ \mathsf{sdec}(x, \mathsf{senc}(x,y)) \to y\}.$$

Our signature and associated rewrite rules allow us to model: asymmetric encryption, pairing, and symmetric encryption.

We define the participants, Alice and Bob, in our composition of the Needham-Schroeder protocol with a symmetric encryption scheme (Figure 4). Hence, the complete composition is modelled by the configuration $E_{\mathsf{NSL}}, \mathcal{P}_{\mathsf{NSL}} = \{c, k_A, k_B\}$, $\{!\nu\, s.A, !B, !\overline{c}\langle\mathsf{pk}(k_A)\rangle, !\overline{c}\langle\mathsf{pk}(k_B)\rangle\}$ (which are implicitly assumed to be running in phase 0).

---

[3] This example was inspired by discussion with Blanchet & Cortier [14]

Fig. 4 Processes modelling Alice and Bob in the Needham-Schroeder protocol

$$
\begin{aligned}
A = &\ \nu\, n_A. \\
&\ \overline{c}\langle\mathsf{aenc}(\mathsf{pk}(k_B), n_A)\rangle. \\
&\ c(x_{\mathsf{ciph}}). \\
&\ \mathsf{let}\ x_{\mathsf{pair}} = \mathsf{adec}(k_A, x_{\mathsf{ciph}})\ \mathsf{in} \\
&\ \mathsf{let}\ y_{\mathsf{nonce}} = \mathsf{fst}(x_{\mathsf{pair}})\ \mathsf{in} \\
&\ \mathsf{if}\ y_{\mathsf{nonce}} = n_A\ \mathsf{then} \\
&\ \mathsf{let}\ x_{\mathsf{nonce}} = \mathsf{snd}(x_{\mathsf{pair}})\ \mathsf{in} \\
&\ \overline{c}\langle\mathsf{aenc}(\mathsf{pk}(k_B), x_{\mathsf{nonce}})\rangle. \\
&\ (\ast\ \mathsf{end\ key\ exchange}\ \ast) \\
&\ \overline{c}\langle\mathsf{pair}(n_A, \mathsf{senc}(x_{\mathsf{nonce}}, s))\rangle
\end{aligned}
$$

$$
\begin{aligned}
B = &\ c(y_{\mathsf{ciph}}). \\
&\ \mathsf{let}\ y_{\mathsf{nonce}} = \mathsf{adec}(k_B, y_{\mathsf{ciph}})\ \mathsf{in} \\
&\ \nu\, n_B. \\
&\ \overline{c}\langle\mathsf{aenc}(\mathsf{pk}(k_A), \mathsf{pair}(y_{\mathsf{nonce}}, n_B))\rangle. \\
&\ c(y'_{\mathsf{ciph}}). \\
&\ \mathsf{let}\ x_{\mathsf{nonce}} = \mathsf{adec}(k_B, y'_{\mathsf{ciph}})\ \mathsf{in} \\
&\ \mathsf{if}\ x_{\mathsf{nonce}} = n_B\ \mathsf{then} \\
&\ c(y_{\mathsf{pair}}). \\
&\ \mathsf{let}\ y'_{\mathsf{nonce}} = \mathsf{fst}(y_{\mathsf{pair}})\ \mathsf{in} \\
&\ \mathsf{if}\ y_{\mathsf{nonce}} = y'_{\mathsf{nonce}}\ \mathsf{then} \\
&\ \mathsf{let}\ y''_{\mathsf{ciph}} = \mathsf{snd}(y_{\mathsf{pair}})\ \mathsf{in} \\
&\ \mathsf{let}\ y_{\mathsf{secret}} = \mathsf{sdec}(n_B, y''_{\mathsf{ciph}})\ \mathsf{in}\ 0
\end{aligned}
$$

5.2 Analysis

We would like to analyse $\widehat{\mathsf{secrecy}}(\{c\}, s) : E_{\mathsf{NSL}}, \mathcal{P}_{\mathsf{NSL}}$ using ProVerif. However, the Horn clauses generated by Blanchet's translation of $\mathcal{P}_{\mathsf{NSL}}$ result in Horn clauses which model the process $\mathcal{P}_{\mathsf{NSL}}$ with a replication after the name restriction $\nu\, n_A$. Unfortunately, secrecy does not hold in this process, because the attacker can learn n_A during one run of the protocol and simulate Bob during a second run to learn Alice's secret s. Indeed, ProVerif finds such an attack. However, this is a false attack, because the protocol states that Alice should use a fresh nonce for every session of the protocol. This false attack can be avoided using our results.

By Corollary 1, to prove $\widehat{\mathsf{secrecy}}(\{c\}, s) : E_{\mathsf{NSL}}, \mathcal{P}_{\mathsf{NSL}}$ it is sufficient to prove $\mathsf{secrecy}(Init, s') : E \cup \{s'\}, \mathcal{Q}$, where $\mathcal{Q} \in \Delta(\{c\}, \mathcal{P}_{\mathsf{NSL}} \cup \{A\{s'/s\}\})$. Let us construct a suitable $\mathcal{Q}$. By definition of Δ (Definition 4), we have $\mathcal{Q} = \{L\} \cup \bigcup_{P \in \mathcal{P}_{\mathsf{NSL}}}\{!R \mid R \in \delta(Init, P)\}$ for some $L \in \delta(\emptyset, A\{s'/s\})$. Since $\mathcal{P}_{\mathsf{NSL}}$ does not contain any private channels, for all contexts C, processes L and terms M and N such that $\mathsf{fv}(M) \cap \{c\} = \emptyset$, we have $\nu\, s.A \neq C[\overline{M}\langle N\rangle.L']$, $B \neq C[\overline{M}\langle N\rangle.L']$, $\overline{c}\langle\mathsf{pk}(k_A)\rangle \neq C[\overline{M}\langle N\rangle.L']$, and $\overline{c}\langle\mathsf{pk}(k_B)\rangle \neq C[\overline{M}\langle N\rangle.L']$. Moreover, since $\overline{c}\langle\mathsf{pk}(k_A)\rangle$ and $\overline{c}\langle\mathsf{pk}(k_B)\rangle$ do not contain inputs, we have $\{!R \mid R \in \delta(Init, \overline{c}\langle\mathsf{pk}(k_A)\rangle)\} = \{!\overline{c}\langle\mathsf{pk}(k_A)\rangle\}$ and, similarly, $\{!R \mid R \in \delta(Init, \overline{c}\langle\mathsf{pk}(k_B)\rangle)\} = \{!\overline{c}\langle\mathsf{pk}(k_B)\rangle\}$. The set $\{!R \mid R \in \delta(Init, \nu\, s.A)\}$ contains $!\nu\, s.A$ and a modified version of $!\nu\, s.A$, namely, $!\nu\, s.A$ with a phase inserted before the input. The set $\{!R \mid R \in \delta(Init, B)\}$ contains $!B$ and three modified versions of $!B$, namely, $!B$ with a phase before each of the three inputs. We have a concrete definition of $\bigcup_{P \in \mathcal{P}_{\mathsf{NSL}}}\{!R \mid R \in \delta(Init, P)\}$ and we let $L \in \delta(\emptyset, A\{s'/s\})$ be $A\{s'/s\}$ with a phase inserted before the penultimate output. ProVerif can automatically verify that $\mathsf{secrecy}(Init, s') : E \cup \{s'\}, \mathcal{Q}$ holds (verification takes less than one second using ProVerif 1.86pl4 on Ubuntu 12.04.3 with 3.60GHz Intel Xeon E5-1620 and 8GB memory).

6   Case study II: Bluetooth Simple Pairing

The Bluetooth Simple Pairing protocol [22, 23] extends the (elliptic curve) Diffie-Hellman protocol to provide authenticated key exchange. There are a number of variations of this protocol, depending on the capabilities of the devices being paired, here we look at the "Numeric Comparison Protocol" that aims to securely pair devices that are capable of displaying a short number on a screen, and receiving an input from the user. The displays of the participants' devices provide a low-bandwidth, authenticated "out of band" channel, which is assumed to be untappable by the attacker. The bandwidth constraint of this out of band channel makes BlueTooth Pairing non-trivial.

The protocol (Figure 5) proceeds as follows. Alice and Bob establish a Diffie-Hellman key (Steps 1 & 2). Bob generates a nonce and outputs a commitment to his nonce (Step 3), Alice outputs a nonce (Step 4), and Bob reveals his nonce (Step 5). Alice and Bob each establish their transcript of values g^x, g^y, N_a, N_b and check that the first few characters of their transcripts – which we write as $\mathtt{short}(g^x, g^y, N_a, N_b)$ – match using their out of band channel (Step 6). Alice and Bob can use the Diffie-Hellman key g^{xy}.

The rationale behind the protocol is that: checking the transcripts on an authenticated channel guarantees the key is shared between Alice and Bob. However, Alice and Bob are only required to check a few characters of the transcript, because the authenticated channel is low-bandwidth, which limits the number of characters that can be checked. Formally, we derive the low-entropy string of the characters that should be checked by applying function $\mathtt{short}$ to the values g^x, g^y, N_a, N_b.

Fig. 5 BlueTooth Simple Pairing

Alice   Bob

1. $\xrightarrow{\quad g^x \quad}$

2. $\xleftarrow{\quad g^y \quad}$

3. $\xleftarrow{\quad \mathtt{H}(g^y, g^x, N_b) \quad}$

4. $\xrightarrow{\quad N_a \quad}$

5. $\xleftarrow{\quad N_b \quad}$

6. $\xleftrightarrow[Out\ of\ Band]{\quad transcript \quad}$

Collision attacks. String $\mathtt{short}(g^x, g^y, N_a, N_b)$ is low-entropy, therefore, unlike computing a hash on g^x, g^y, N_a, N_b, it is computationally feasible to find a collision. The following example demonstrates collision attacks.

Example 4. Consider a variant of BlueTooth Simple Pairing, without step 3. In this protocol, the attacker can replace Alice's message g^x with g^z and Bob's message N_b with N_e such that $\mathtt{short}(g^z, g^y, N_a, N_b) = \mathtt{short}(g^x, g^z, N_a, N_e)$. (This would require the attacker to make many guesses for the nonce N_e.) Given that Alice and Bob would now see the same low-entropy string, the attacker can impersonate Alice.

Such an attack is prevented in BlueTooth Simple Pairing, by forcing both participants to commit to their nonces before they see the other participant's nonce. Alice must send her nonce first, and Bob must send a hash of this nonce before he sees Alice's Nonce. Alice will check that this nonce matches the hash

between steps 5 and 6 of the protocol. This ensures that the attacker never has the opportunity to launch an impersonation attack.

We remark that collision attacks are conceptually different from guessing attacks against weak secrets [6, 16], which occur when an attacker gets enough information from a protocol to verify a guess of a value, which is low entropy. The kind of brute force attack we describe here occurs when an attacker has some control over the inputs to a high-entropy function and needs to force the short code based on the output of that function to equal a certain value. In particular, if an attacker controls one input and knows the values of all the other inputs, the attacker can generated a large number of possible inputs and, by brute force, find an input of their own that makes the output of the function produce a value that matches any short code they wish.

6.1   Applied pi calculus model

Our model assumes that an attacker can generate sufficiently many transcripts such that distinct transcripts share the same first few characters. That is, given $\texttt{short}(M_1, M_2, M_3, M_4)$, the adversary can compute $\texttt{short}(N_1, N_2, N_3, N_4)$ such that $\texttt{short}(M_1, M_2, M_3, M_4) = \texttt{short}(N_1, N_2, N_3, N_4)$. Function $\texttt{bruteforce}$ captures this. Moreover, we supplement $\texttt{short}$ with $\texttt{shortb}$ to enable automated analysis.

We construct a signature $\Sigma = \{\texttt{g}, \texttt{f}, \texttt{sdec}, \texttt{senc}, \texttt{H}, \texttt{bruteforce}, \texttt{short}, \texttt{shortb}\}$, where $\texttt{g}$ is a unary function, $\texttt{sdec}$, $\texttt{senc}$ and $\texttt{f}$ are binary functions, $\texttt{H}$ is a ternary function, and $\texttt{bruteforce}$, $\texttt{short}$ and $\texttt{shortb}$ take four arguments. Function $\texttt{H}$ represents a hash function, $\texttt{sdec}$ and $\texttt{senc}$ capture symmetric encryption, i.e., $\texttt{def}(\texttt{short}) = \{\texttt{sdec}(x, \texttt{senc}(x, y)) \to y\}$, and the purpose of the remaining functions is explained below.

Diffie-Hellman key agreement is modelled in the standard fashion [10, §9.1] using functions $\texttt{f}$ and $\texttt{g}$ and the following equation:

$$\texttt{f}(x, \texttt{g}(y)) = \texttt{f}(y, \texttt{g}(x))$$

which allows us to capture $(g^x)^y = (g^y)^x$. (A more general setting is beyond the scope of [10, §9.1].)

Function $\texttt{shortb}$ is used to model a low-entropy string derived from its input. For instance, $\texttt{shortb}(g^x, g^y, N_A, N_B)$ represents a low-entropy string derived from terms g^x, g^y, N_A, and N_B. To capture collision attacks, function $\texttt{shortb}$ is not used directly (that is, we do not use $\texttt{shortb}$ in processes), instead, destructor $\texttt{short}$ is used and we map occurrences of $\texttt{short}$ to $\texttt{shortb}$ using the following rewrite rule:

$$\texttt{short}(w, x, y, z) \to \texttt{shortb}(w, x, y, z)$$

Moreover, we add the following rewrite rules to capture collisions:

$$\texttt{short}(w, x, y, \texttt{bruteforce}(w, x, y, \texttt{shortb}(\hat{w}, \hat{x}, \hat{y}, \hat{z}))) \to \texttt{shortb}(\hat{w}, \hat{x}, \hat{y}, \hat{z})$$
$$\texttt{short}(w, x, \texttt{bruteforce}(w, x, z, \texttt{shortb}(\hat{w}, \hat{x}, \hat{y}, \hat{z})), z) \to \texttt{shortb}(\hat{w}, \hat{x}, \hat{y}, \hat{z})$$
$$\texttt{short}(w, \texttt{bruteforce}(w, y, z, \texttt{shortb}(\hat{w}, \hat{x}, \hat{y}, \hat{z})), y, z) \to \texttt{shortb}(\hat{w}, \hat{x}, \hat{y}, \hat{z})$$
$$\texttt{short}(\texttt{bruteforce}(x, y, z, \texttt{shortb}(\hat{w}, \hat{x}, \hat{y}, \hat{z})), x, y, z) \to \texttt{shortb}(\hat{w}, \hat{x}, \hat{y}, \hat{z})$$

It follows from our rewrite rules that

$$\texttt{short}(M_1, M_2, M_3, M_4) \rightarrow \texttt{shortb}(M_1, M_2, M_3, M_4)$$

and

$$\texttt{short}(N_1, N_2, N_3, \texttt{bruteforce}(N_1, N_2, N_3, \texttt{shortb}(M_1, M_2, M_3, M_4)))$$
$$\rightarrow \texttt{shortb}(M_1, M_2, M_3, M_4))$$

i.e., terms $\texttt{short}(M_1, M_2, M_3, M_4)$ and $\texttt{short}(N_1, N_2, N_3, \texttt{bruteforce}(N_1, N_2, N_3,$ $\texttt{shortb}(M_1, M_2, M_3, M_4)))$ collide.

The ProVerif source of this, and all our other examples, are available from the following URL: `http://www.cs.bham.ac.uk/~tpc/projects/falseattacks`.

6.2 Analysis

We analyse this protocol by adding in the exchange of a value at the end of the protocol using the key g^{xy} and testing for the secrecy of this value. Running our model in ProVerif we find that it results in a false attack, due to the commitment problem that was discussed above. The tool suggests that the attacker can send a dummy commitment value and then observe nonces N_a and N_b. The attacker can then use a brute forced value based on N_b and start the protocol again with a commitment to this new brute forced value, however, in this second run of the protocol Bob would be using a different N_b value and so the attack is a false one.

Applying Theorem 1 to transform the BlueTooth Protocol results in a model (available from the aforementioned URL) for which ProVerif can verify that the secrecy of *seckey* in less than a second, and so Theorem 1 then tells us that, in spite of the false attack, the original protocol is also secure. If we remove the check of the commitment sent in step 3 of the protocol, then ProVerif finds the attack in which the attacker performs a man in the middle attack using the `bruteforce` function to find a value which matches the short code sent between Alice and Bob.

6.3 Case study III: the Simplified Simple Pairing Protocol

Our final case study proposes a simplified version of the Bluetooth Simple Pairing protocol (Figure 6). Our new protocol merges the tasks of the Diffie-Hellman exponents and the nonces. The out of band channel is used for confirmations of a short code based on the Diffie-Hellman key. We note that in the specification of the Bluetooth Simple Pairing protocol [23] the Diffie-Hellman steps of the protocol are presented as a method of stopping eavesdropping attacks, and the steps using the nonce values are presented quite separately as a method of stopping active attacks; the contribution of our protocol is to show that these steps can be merged. We also note that in the Bluetooth Simple Pairing protocol the channel is protected by the entropy of the secret Diffie-Hellman values and the

nonces, whereas in our new protocol it is just protected by the entropy of the secret Diffie-Hellman values. Therefore, for our protocol to be secure, the secret Diffie-Hellman values must be high-entropy and fresh; this is not currently an explicit requirement stated in the Bluetooth specification.

Fig. 6 SSP: Simplified Simple Pairing

Alice   Bob

1. $\xrightarrow{\quad Hash(g^x) \quad}$

2. $\xleftarrow{\quad g^y \quad}$

3. $\xrightarrow{\quad g^x \quad}$

Verify $3.g^x$ vs. $1.\text{Hash}(g^x)$

4. $\overset{short(g^{xy})}{\underset{Out\ of\ Band}{\rightleftharpoons}}$

Verify $4.\text{short}(g^{xy})$ vs. g^{xy} Verify $4.\text{short}(g^{xy})$ vs. g^{xy}

5. $\xrightarrow{\quad \{seckey\}_{g^{xy}} \quad}$

Our new protocol starts with Alice sending a commitment to a particular Diffie-Hellman exponent. Bob then sends his exponent and Alice replies by sending her's. Bob must check that the exponent he receives matches the commitment to avoid the brute force attack described above. Alice and Bob then compare the transcript of the key on their out of bound channel. The nonces used in the Bluetooth Simple Pairing protocol are no longer needed, as the freshness guarantees are now provided by the values x and y.

A brute force attack may be tried against the short string in this protocol. To allow this in our model, we adapt the rules for `bruteforce` and `short` to work against Diffie-Hellman exponents. In this case, the `short` function is applied to the `f` function i.e., `short(f(s,g(t)))` represents the short string generated from the Diffie-Hellman key `f(s,g(t))`.

Given `g(t)` and `M`, the attacker can derive a term `N` such that `short(f(N,g(t)))` `= short(M)`, and we write this `N` value as `bruteforce(g(t),M)` and capture the desired relation using the following rewrite rules:

$$\text{short}(x) \rightarrow \text{shortb}(x)$$
$$\text{short}(\text{f}(\text{bruteforce}(x, \text{shortb}(y)), x)) \rightarrow \text{shortb}(y)$$
$$\text{short}(y, \text{f}(\text{bruteforce}(y, \text{shortb}(x)))) \rightarrow \text{shortb}(x)$$

It follows from the above rules that given `g(t)` and `f(s,g(t))`, the attacker can derive `bruteforce(g(t),short(f(s,g(t))))` such that `short(f(bruteforce(g(t),` `short(f(s,g(t)))),g(t))) = short(f(s,g(t)))`, thereby modelling a brute force attack against the Diffie-Hellman exponents.

Analysing this scheme using ProVerif results in the discovery of a false attack, whereas Theorem 1 allows us to verify the secrecy of *seckey*.

7   Conclusion

We have shown how false atacks can be avoided when analysing secrecy with ProVerif. Our method works by inserting phases into processes such that they enforce an ordering on Horn clauses. We demonstrate the applicability of our methodology by analysing BlueTooth Pairing. This case study leads use to develop theory to enable the analysis of protocols that are vulnerable to impersonation attacks. Finally, we show that BlueTooth Pairing can be simplified and show that the simplified scheme satisfies the same secrecy objectives.

References

1. Martín Abadi, Bruno Blanchet, and Cédric Fournet. Just Fast Keying in the Pi Calculus. *ACM Transactions on Information and System Security*, 10(3), 2007.
2. Martín Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In *POPL'01: 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 104–115. ACM Press, 2001.
3. Xavier Allamigeon and Bruno Blanchet. Reconstruction of Attacks against Cryptographic Protocols. In *CSFW'05: 18th Computer Security Foundations Workshop*, pages 140–154. IEEE Computer Society, 2005.
4. Myrto Arapinis, Véronique Cortier, Steve Kremer, and Mark Ryan. Practical everlasting privacy. In *POST'13: 2nd International Conference on Principles of Security and Trust*, volume 7796 of *LNCS*, pages 21–40. Springer, 2013.
5. Michael Backes, Cătălin Hrițcu, and Matteo Maffei. Automated Verification of Remote Electronic Voting Protocols in the Applied Pi-calculus. In *CSF'08: 21st IEEE Computer Security Foundations Symposium*, pages 195–209. IEEE Computer Society, 2008.
6. Mathieu Baudet. Deciding security of protocols against off-line guessing attacks. In *In Proc. 12th ACM Conference on Computer and Communications Security (CCS?05*, pages 16–25. ACM Press, 2005.
7. Mathieu Baudet. *Sécurité des protocoles cryptographiques : aspects logiques et calculatoires*. PhD thesis, Laboratoire Spécification et Vérification, ENS Cachan, France, 2007.
8. Bruno Blanchet. An efficient cryptographic protocol verifier based on Prolog rules. In *CSFW'01: 14th IEEE Computer Security Foundations Workshop*, pages 82–96. IEEE Computer Society, 2001.
9. Bruno Blanchet. From Secrecy to Authenticity in Security Protocols. In *SAS'02: 9th International Static Analysis Symposium*, volume 2477 of *LNCS*, pages 342–359. Springer, 2002.
10. Bruno Blanchet. Automatic Verification of Correspondences for Security Protocols. *Journal of Computer Security*, 17(4):363–434, 2009.
11. Bruno Blanchet. Private email communication, 12th November 2012.

12. Bruno Blanchet. Security Protocol Verification: Symbolic and Computational Models. In *POST'12: 1st Conference on Principles of Security and Trust*, volume 7215 of *LNCS*, pages 3–29. Springer, 2012.

13. Bruno Blanchet, Martín Abadi, and Cédric Fournet. Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming*, 75(1):3–51, February–March 2008.

14. Bruno Blanchet and Véronique Cortier. Private email communication, 13 November 2012.

15. Bruno Blanchet and Ben Smyth. ProVerif: Automatic Cryptographic Protocol Verifier User Manual & Tutorial. `http://www.proverif.ens.fr/`, 2011.

16. Bruno Blanchet, Ben Smyth, and Vincent Cheval. Proverif 1.88: Automatic cryptographic protocol verifier, user manual and tutorial, 2013.

17. Richard Chang and Vitaly Shmatikov. Formal analysis of authentication in bluetooth device pairing. In *Foundations of Computer Security and Automated Reasoning for Security Protocol Analysis*, 2007.

18. Liqun Chen and Mark D. Ryan. Attack, Solution and Verification for Shared Authorisation Data in TCG TPM. In *FAST'09: 6th International Workshop on Formal Aspects in Security and Trust*, volume 5983 of *LNCS*, pages 201–216. Springer, 2009.

19. Stéphanie Delaune, Steve Kremer, Mark Dermot Ryan, and Graham Steel. Formal analysis of protocols based on tpm state registers. In *CSF'11: 24th IEEE Computer Security Foundations Symposium*, pages 66–80. IEEE, 2011.

20. Stéphanie Delaune, Mark D. Ryan, and Ben Smyth. Automatic verification of privacy properties in the applied pi-calculus. In *IFIPTM'08: 2nd Joint iTrust and PST Conferences on Privacy, Trust Management and Security*, volume 263 of *International Federation for Information Processing (IFIP)*, pages 263–278. Springer, 2008.

21. Nancy A. Durgin, Patrick Lincoln, John C. Mitchell, and A. Scedrov. Multiset rewriting and the complexity of bounded security protocols. *Journal of Computer Security*, 12(2):247–311, 2004.

22. Bluetooth Special Interest Group. Specification of the bluetooth system, 2001.

23. Bluetooth Special Interest Group. Simple pairing whitepaper, 2006.

24. Gavin Lowe. Breaking and Fixing the Needham-Schroeder Public-Key Protocol using FDR. In *TACAS'96: 2nd International Workshop on Tools and Algorithms for the Construction and Analysis of Systems*, volume 1055 of *LNCS*, pages 147–166. Springer, 1996.

25. Catherine Meadows. Open Issues in Formal Methods for Cryptographic Protocol Analysis. In *MMM-ACNS'01: International Workshop on Information Assurance in Computer Networks*, volume 2052 of *LNCS*, page 21. Springer, 2001.

26. Roger M. Needham and Michael D. Schroeder. Using Encryption for Authentication in Large Networks of Computers. *Communications of the ACM*, 21(12):993–999, 1978.

27. Mark D. Ryan and Ben Smyth. Applied pi calculus. In Véronique Cortier and Steve Kremer, editors, *Formal Models and Techniques for Analyzing Security Protocols*, chapter 6. IOS Press, 2011.

28. Ben Smyth, Mark D. Ryan, and Liqun Chen. Formal analysis of privacy in Direct Anonymous Attestation schemes, 2012.

29. Fangming Zhao, Yoshikazu Hanatani, Yuichi Komano, Ben Smyth, Satoshi Ito, and Toru Kambayashi. Secure Authenticated Key Exchange with Revocation for Smart Grid. In *ISGT'12: 3rd IEEE Power & Energy Society Conference on Innovative Smart Grid Technologies*, pages 1–8, 2012.