

Applied pi calculus*

Mark D. Ryan and Ben Smyth

School of Computer Science,
University of Birmingham,
United Kingdom

July 2010

Revised: November 2010, March 2011,
November 2012, March 2013.

Abstract

The applied pi calculus is a language for modelling security protocols. It is an extension of the pi calculus, a language for studying concurrency and process interaction. This chapter presents the applied pi calculus in a tutorial style. It describes reachability, correspondence, and observational equivalence properties, with examples showing how to model secrecy, authentication, and privacy aspects of protocols.

Keywords: Cryptographic protocols, protocol verification, formal methods, reachability, correspondence properties, observational equivalence, tutorial.

1 Introduction

The applied pi calculus [AF01] is a language for describing and analysing security protocols. It provides an intuitive process syntax for detailing the actions of the participants in a protocol, emphasising their communication. The syntax is coupled with a formal semantics to allow reasoning about protocols. The language is based on the pi calculus with the addition of a rich term algebra to enable modelling of the cryptographic operations used by security protocols. A wide variety of cryptographic primitives can be abstractly modelled by means of an equational theory. The calculus allows one to express several types of security goal, and to analyse whether the protocol meets its goal or not. This analysis can sometimes be performed automatically, using the ProVerif software tool [BAF08, Bla09, BS10].

The applied pi calculus has been used to model security protocols in a variety of areas. The following examples are not exhaustive:

- Certified email [AB05b];
- Privacy properties [KR05, DKR09b, BHM08, KT09], and election verifiability properties [KRS10, SRKK10] in electronic voting;
- Authorisation protocols [CR09, MGR09], and attestation protocols [SRC07, BMU08] in trusted computing;
- Interoperability of web services, where a compiler [BFGT06] converts descriptions in the F# language [SGC07] into applied pi, suitable for analysis by ProVerif;


*This work is an extended version of Mark D. Ryan and Ben Smyth (2011) *Applied pi calculus*. In V Cortier and S Kremer, editor, *Formal Models and Techniques for Analyzing Security Protocols*.

- Integrity of file systems on untrusted storage [BC08];
- Authentication protocols and key agreement [ABF07].

Notation and running example. Throughout this Chapter we will make use of the Handshake protocol discussed in Section 1.1. The following notation will be used, symmetric encryption of a message m with key k is denoted $\{m\}_k^s$. Similarly, we write $\{m\}_{pk}^a$ for the asymmetric encryption of m with the public key pk and $[m]_{sk}$ for the digital signature of m with the secret key sk . The pairing of messages m_1 and m_2 is denoted $\langle m_1, m_2 \rangle$.

1.1 Handshake protocol

A naïve Handshake protocol between a client \mathcal{C} and server \mathcal{S} is illustrated below. It is assumed that each of them has a public/private key pair, and that the client knows the server’s public key $pk(sk_{\mathcal{S}})$. The aim of the protocol is to establish a secret symmetric key k , enabling the client to communicate a secret s to the server. The protocol proceeds as follows. On request from a client \mathcal{C} , server \mathcal{S} generates a fresh session key k , signs it with her private key $sk_{\mathcal{S}}$ and encrypts it using her client’s public key $pk(sk_{\mathcal{C}})$. When \mathcal{C} receives this message he decrypts it using his private key $sk_{\mathcal{C}}$, verifies the digital signature made by \mathcal{S} using her public key $pk(sk_{\mathcal{S}})$, and extracts the session key k . \mathcal{C} uses this key to symmetrically encrypt the secret s and sends the encrypted message to \mathcal{S} . The rationale behind the protocol is that \mathcal{C} receives the signature asymmetrically encrypted with his public key and hence he should be the only one able to decrypt its content. Moreover, the digital signature should ensure that \mathcal{S} is the originator of the message. The protocol narration is illustrated as follows:



Note that protocol narrations (as above) are useful, but lack detail. For example, they do not specify the construction of nonces, nor do they describe any checks which should be made by the participants during the execution of the protocol. Such checks include verifying digital signatures and ensuring that encrypted messages are correctly formed. Failure of these checks typically results in the participant aborting the protocol. These details will be explicitly stated when protocols are modelled in the applied pi calculus. (For further discussion on protocol specification see [AN96, Aba00].)

Informally, the three properties we would like this protocol to provide are:

1. Secrecy: The value s is known only to \mathcal{C} and \mathcal{S} .
2. Authentication of \mathcal{S} : if \mathcal{C} reaches the end of the protocol with session key k , then \mathcal{S} proposed k for use by \mathcal{C} .
3. Authentication of \mathcal{C} : if \mathcal{S} reaches the end of the protocol and believes session key k has been shared with \mathcal{C} , then \mathcal{C} was indeed her interlocutor and has k .

The different forms of the two authentication properties arise because of the different assumptions we made about \mathcal{C} and \mathcal{S} . Recall that \mathcal{C} knows \mathcal{S} ’s public key, and is only willing to run the protocol with \mathcal{S} . But \mathcal{S} is willing to run the protocol with anyone.

Careful analysis reveals that the protocol does not satisfy all three of the intended properties. It is vulnerable to a *man-in-the-middle* attack (illustrated below). If a dishonest participant \mathcal{M} starts a session with \mathcal{S} , then \mathcal{M} is able to impersonate \mathcal{S} in a subsequent session he starts with \mathcal{C} . At the end of the protocol \mathcal{C} believes that he shares the secret s with \mathcal{S} , while he actually shares s with \mathcal{M} .

The grammar for *processes* is shown below:

$P, Q, R ::=$	processes (or plain processes)
0	null process
$P \mid Q$	parallel composition
$!P$	replication
$\nu n.P$	name restriction
if $M = N$ then P else Q	conditional
$u(x).P$	message input
$\bar{u}\langle M \rangle.P$	message output

The null process 0 does nothing; $P \mid Q$ is the parallel composition of processes P and Q , used to represent participants of a protocol running in parallel; and replication $!P$ is the infinite composition $P \mid P \mid \dots$, which is often used to capture an unbounded number of sessions. Name restriction $\nu n.P$ binds n inside P , the introduction of restricted names (or private names) is useful to capture both fresh random numbers (modelling nonces and keys, for example) and private channels. The conditional if $M = N$ then P else Q is standard, but we stress $M = N$ represents equality (modulo an equational theory) rather than strict syntactic identity. For convenience we abbreviate conditionals as if $M = N$ then P , when Q is the null process. Finally, communication is captured by message input and message output. The process $u(x).P$ awaits a message from channel u and then behaves as P with the received message bound to the variable x ; that is, every free occurrence of x in P refers to the message received. The process $\bar{u}\langle M \rangle.P$ is ready to send M on channel u and then run P . In both of these cases we may omit P when it is 0 . We write $P\{M/x\}$ for P with all free occurrences of x replaced by M and sometimes we write let $x = M$ in P instead of $P\{M/x\}$; in such substitutions, we insist that no name or variable occurring in M becomes bound by a restriction occurring in P (for example, let $x = c$ in $\nu c.\bar{c}\langle x \rangle$ is not allowed, but let $x = c$ in $\nu a.\bar{a}\langle x \rangle$ is permitted).

Bracketing must be used to avoid ambiguities in the way processes are written down. For example, the process $!P \mid Q$ might be interpreted as $(!P) \mid Q$ or as $!(P \mid Q)$. These processes are different. To avoid too much bracketing, we adopt conventions about the precedence of process operators. Unary operators $!$, νn , $u(x)$, and $\bar{u}\langle M \rangle$ bind more closely than binary operators; and the binary if-then-else operator binds more closely than the binary operator \mid . It follows that the expression $c(x).\text{if } x = M \text{ then } P \mid !Q \mid R$ means $(c(x).\text{if } x = M \text{ then } P) \mid (!Q) \mid R$. We remark that different conventions are used elsewhere, for example [BS10].

The expression $P \mid Q \mid R$ is also ambiguous, since it could mean either $(P \mid Q) \mid R$ or $P \mid (Q \mid R)$. However, we will later see that these processes are semantically identical, so we tolerate the ambiguity in the syntax. Another possible ambiguity arises because of the convention of omitting “else 0” in the if-then-else construct: it is not clear which “if” the “else” applies to in the expression:

$$\text{if } M = N \text{ then if } K = L \text{ then } Q \text{ else } R.$$

In absence of brackets indicating the contrary, we adopt the convention that the else branch belongs to the closest if and hence the statement should be interpreted as if $M = N$ then (if $K = L$ then Q else R).

Processes are extended with *active substitutions* to capture the knowledge exposed to the adversarial environment:

$A, B, C ::=$	extended processes
P	plain process
$A \mid B$	parallel composition
$\nu n.A$	name restriction
$\nu x.A$	variable restriction
$\{M/x\}$	active substitution

The active substitution $\{M/x\}$ represents a process that has previously output M . The value M is now available to the environment by reference to the ‘handle’ x . The active substitution $\{M/x\}$

can replace the variable x for the term M in every process it comes into contact with (formally this is captured by the operational semantics). This behaviour can be controlled by restriction, and the process $\nu x.(\{M/x\} \mid P)$ corresponds exactly to let $x = M$ in P . This allows access to terms which the environment cannot construct; such a scenario may arise, for example, when the environment does not know all of the names occurring in a term. Arbitrarily large active substitutions can be obtained by parallel composition and we occasionally abbreviate $\{M_1/x_1\} \mid \dots \mid \{M_l/x_l\}$ as $\{M_1/x_1, \dots, M_l/x_l\}$ or $\{\tilde{M}/\tilde{x}\}$. The letters σ and τ range over substitutions. We write $N\sigma$ for the result of applying σ to the variables of N . Active substitutions are always assumed to be cycle-free: given substitution σ , the repeated application of σ to itself may only result in a bounded number of replacements. For example, the substitution $\{g(y)/x, n/y\}$ is cycle-free; but $\{g(y)/x, x/y\}$ is not and repeated application could result in terms of unbounded length. Extended processes must have at most one active substitution for each variable and there is exactly one when the variable is under restriction. In particular, this means that νx can only occur if there is an active substitution $\{M/x\}$ in its scope. Finally, we write $\nu \tilde{u}$ for the (possibly empty) series of pairwise-distinct binders $\nu u_1.\nu u_2.\dots.\nu u_l$.

The type system for terms is extended to processes. It enforces that M, N are of the same type in the conditional expression if $M = N$ then P else Q ; message input $u(x)$ is defined only where u is of type $\text{Channel}\langle\omega\rangle$ and x is of type ω ; and similarly message output $\bar{u}(M)$ requires u of type $\text{Channel}\langle\omega\rangle$ and M of type ω . In addition, we assume that in active substitutions $\{M/x\}$, the term M and the variable x are of the same base type. This assumption was not explicitly stated in [AF01] but its necessity has been confirmed [AF06] and is essential for the validity of Theorem 1 [BJPV09]. Finally, we assume extended processes are well-typed.

The *scope* of names and variables are delimited by binders $u(x)$ and νu . The set of bound names $\text{bn}(A)$ contains every name n which is under restriction νn inside A . The set of bound variables $\text{bv}(A)$ consists of all those variables x occurring in A that are bound by restriction νx or input $u(x)$. We also define the set of free names and the set of free variables. The set of free names in A , denoted $\text{fn}(A)$, consists of those names n occurring in A not in the scope of the binder νn . The set of free variables $\text{fv}(A)$ contains the variables x occurring in A which are not in the scope of a restriction νx or input $u(x)$. Note that a name or variable can occur both free and bound in A . Thus, the sets $\text{fv}(A) \cap \text{bv}(A)$ and $\text{fn}(A) \cap \text{bn}(A)$ may be non-empty. We occasionally write $\text{fn}(M)$ and $\text{fv}(M)$, for the set of names, respectively variables, which appear in term M . The concept of bound and free values is similar to local and global scope in programming languages. An extended process is *closed* when every variable x is either bound or defined by an active substitution, (that is, the process contains $\{M/x\}$ for some term M). Example 1 demonstrates name and variable scope. Bound names and bound variables are subject to α -conversion (also called α -renaming); that is, they may be uniformly renamed without changing the meaning of the process, as demonstrated in Example 2.

A *frame*, denoted φ or ψ , is an extended process built from 0 and active substitutions $\{M/x\}$; which are composed by parallel composition and restriction. The *domain* $\text{dom}(\varphi)$ of a frame φ is the set of variables that φ exports; that is, the set of variables x for which φ contains an active substitution $\{M/x\}$ such that x is not under restriction. Every extended process A can be mapped to a frame $\varphi(A)$ by replacing every plain process in A with 0 . The frame $\varphi(A)$ represents the static knowledge output by a process to its environment. The domain $\text{dom}(A)$ of A is the domain of $\varphi(A)$.

Example 1 (Name and variable scope) Consider the closed process $A \triangleq \nu x.((c(y).\bar{c}\langle \text{aenc}(\text{pair}(x, z), y) \rangle) \mid (\nu s.\{s/x\} \mid \{h(x, s')/z\}))$ defined with respect to the signature $\Sigma = \{\text{aenc}, h, \text{pair}\}$ where $\text{aenc}, h, \text{pair}$ are all binary functions. The occurrences of the variable x are bound by the variable restriction νx and those of y are bound by the input $c(y)$. The name s is bound by the name restriction νs . The remaining name s' and variable z occur free. To summarise, we have $\text{fv}(A) = \{z\}$, $\text{bv}(A) = \{x, y\}$, $\text{fn}(A) = \{s'\}$ and $\text{bn}(A) = \{s\}$. Now let $A' \triangleq A \mid \{s/y\}$, and observe that the occurrences of s and y in the active substitution $\{s/y\}$ are free. They have no relation to the bound values s, y which occur in A . To avoid confusion it is good practice to use distinct identifiers for free and bound, names and variables.

Example 2 (α -conversion) α -conversion means renaming a bound name or variable without changing the semantics. This is also done in logic (the statements $\forall x.P(x)$ and $\forall y.P(y)$ are equivalent), and in programming (the programs *for i in I do $P(i)$* and *for j in I do $P(j)$* are equivalent). In the applied pi calculus, the process $\nu k.c(x).\bar{c}\langle \text{senc}(k,x) \rangle$ is equivalent to $\nu s.c(y).\bar{c}\langle \text{senc}(s,y) \rangle$; we have α -renamed all occurrences of the name k to the name s and similarly the variable x has been renamed y . But, $\nu k.c(x).\nu k.\bar{c}\langle \text{senc}(k,x) \rangle$ is not α -equivalent to $\nu k.c(y).\nu s.\bar{c}\langle \text{senc}(k,y) \rangle$.

2.1.1 Modelling the Handshake protocol

The Handshake protocol (Section 1.1) is defined with respect to the signature Σ_H , which is used to capture primitives modelling cryptographic operators, data structures and constants: $\Sigma_H = \{\text{true}, \text{fst}, \text{snd}, \text{hash}, \text{pk}, \text{getmsg}, \text{pair}, \text{sdec}, \text{senc}, \text{adec}, \text{aenc}, \text{sign}, \text{checksign}, \text{mac}\}$ where true is a constant; $\text{fst}, \text{snd}, \text{hash}, \text{pk}, \text{getmsg}$ are unary functions; and $\text{pair}, \text{sdec}, \text{senc}, \text{adec}, \text{aenc}, \text{sign}, \text{checksign}, \text{mac}$ are binary functions. The behaviour of these functions is captured by the smallest equational theory E_H satisfying the following equations over variables x, y :

$$\begin{aligned} \text{fst}(\text{pair}(x, y)) &= x \\ \text{snd}(\text{pair}(x, y)) &= y \\ \text{sdec}(x, \text{senc}(x, y)) &= y \\ \text{adec}(x, \text{aenc}(\text{pk}(x), y)) &= y \\ \text{getmsg}(\text{sign}(x, y)) &= y \\ \text{checksign}(\text{pk}(x), \text{sign}(x, y)) &= \text{true} \end{aligned}$$

This theory allows us to model: pairing, both symmetric and asymmetric cryptography, digital signature schemes with message recovery, hash functions and message authentication codes. For example, in order to express that the application of the symmetric decryption function sdec to the term modelling a symmetric encryption $\text{senc}(k, m)$ should return the plaintext m if the correct key k is supplied, we use the equation $\text{sdec}(x, \text{senc}(x, y)) = y$. The absence of any equational theory associated with the hash function ensures preimage resistance, second-preimage resistance and collision resistance properties of cryptographic hash functions (in fact, far stronger properties are ensured); and similar properties for the function mac . Observe that a tuple M_1, \dots, M_l of an arbitrary length l can be constructed by pairing $\text{pair}(M_1, \text{pair}(M_2, \text{pair}(\dots, \text{pair}(M_{l-1}, M_l) \dots)))$ and an element can be extracted using the equations defined for fst , snd .

The Handshake protocol can now be captured in our calculus as the process P , defined as follows.

$$\begin{aligned} P &\triangleq \nu sk_S.\nu sk_C.\nu s. \\ &\quad \text{let } pk_S = \text{pk}(sk_S) \text{ in let } pk_C = \text{pk}(sk_C) \text{ in} \\ &\quad (\bar{c}\langle pk_S \rangle \mid \bar{c}\langle pk_C \rangle \mid !P_S \mid !P_C) \\ P_S &\triangleq c(x.pk).\nu k.\bar{c}\langle \text{aenc}(x.pk, \text{sign}(sk_S, k)) \rangle. \\ &\quad c(z).\text{if } \text{fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\ P_C &\triangleq c(y).\text{let } y' = \text{adec}(sk_C, y) \text{ in let } y.k = \text{getmsg}(y') \text{ in} \\ &\quad \text{if } \text{checksign}(pk_S, y') = \text{true} \text{ then} \\ &\quad \bar{c}\langle \text{senc}(y.k, \text{pair}(\text{tag}, s)) \rangle \end{aligned}$$

The process begins by constructing the private keys sk_C, sk_S for principals \mathcal{C}, \mathcal{S} respectively. The public key parts $\text{pk}(sk_C), \text{pk}(sk_S)$ are then output on the public communication channel c , ensuring they are available to the adversary. (Observe that this is done using handles pk_C and pk_S for convenience.) The protocol then launches multiple copies of processes P_C, P_S representing multiple sessions of the roles of \mathcal{C} and \mathcal{S} . Note that syntactic scope does not represent the knowledge of a protocol's participants. For example, the server's private key sk_S is assumed not to be known by the client \mathcal{C} (hence it does not occur in P_C), even though sk_S is in the scope of P_C .

Figure 1 Properties of equality modulo the equational theory

Given an equational theory E the following properties are satisfied for all terms L, M, N , functions f of arity l , substitutions σ and names m, n :

1. $M = N \in E \Rightarrow M =_E N$
2. Equivalence relation:
 - Reflexivity: $M =_E M$
 - Symmetry: $M =_E N \Rightarrow N =_E M$
 - Transitivity: $L =_E M \wedge M =_E N \Rightarrow L =_E N$
3. Application of function symbols: $M_1 =_E N_1 \wedge \dots \wedge M_l =_E N_l \Rightarrow f(M_1, \dots, M_l) =_E f(N_1, \dots, N_l)$
4. Substitution of terms for variables: $M =_E N \Rightarrow M\sigma =_E N\sigma$
5. Bijective renaming of names: $M =_E N \Rightarrow M\{m/n\} =_E N\{m/n\}$, where $M\{m/n\}$ (respectively $N\{m/n\}$) is the term M (respectively N) with every name n replaced by m .

We assume that \mathcal{S} is willing to run the protocol with any other principal; the choice of her interlocutor will be made by the environment. This is captured by modelling the first input $c(x.pk)$ to $P_{\mathcal{S}}$ as the interlocutor's public key. \mathcal{C} on the other hand only wishes to share his secret s with \mathcal{S} , and \mathcal{C} is assumed to know \mathcal{S} 's public key; accordingly, \mathcal{S} 's public key is hard-coded into the process $P_{\mathcal{C}}$. We additionally assume that each principal is willing to engage in an unbounded number of sessions and hence $P_{\mathcal{C}}$, $P_{\mathcal{S}}$ are under replication.

On request from her interlocutor, server \mathcal{S} starts the protocol by selecting key k and outputting $\text{aenc}(x.pk, \text{sign}(sk_{\mathcal{S}}, k))$; that is, her signature on the key k encrypted with her interlocutor's public key $x.pk$. Meanwhile \mathcal{C} awaits the input of his interlocutor's signature on the key k encrypted using his public key. \mathcal{C} decrypts the message and verifies the signature. Next, if \mathcal{C} believes he is indeed talking to \mathcal{S} , he outputs his secret s encrypted with the symmetric key k . Note that he inserts a tag (modelled as a free name), so that the decryptor can verify that the decryption has worked correctly. Principal \mathcal{S} inputs z and confirms the presence of the tag. Finally, principal \mathcal{S} executes the process Q . The description of Q is independent of the protocol, but one would expect the recovery of the interlocutor's secret; that is, Q is defined by the process $\text{let } z.s = \text{snd}(\text{sdec}(k, z))$ in Q' for some Q' .

The purpose of the protocol is to establish a session key to transport the secret represented by s . We abstract away from the details of what s is, and how many such secrets there are, by modelling it simply as a restricted name.

2.2 Operational semantics

The signature Σ is equipped with an equational theory E , that is, a set of equations of the form $M = N$, where terms M, N are defined over the signature Σ (sometimes written $M, N \in T_{\Sigma}$). This allows us to capture relationships between primitives defined in Σ . We define equality modulo the equational theory, written $=_E$, as the smallest equivalence relation on terms, that contains E and is closed under application of function symbols, substitution of terms for variables and bijective renaming of names. The properties of $=_E$ are summarised in Figure 1. We write $M =_E N$ when the equation $M = N$ is in the theory E and keep the signature implicit. When E is clear from the context we may abbreviate $M =_E N$ as $M = N$. The negation of $M =_E N$ is denoted $M \neq_E N$ (and similarly abbreviated $M \neq N$). For further discussion on equational theories see [AF01, §3] and [AC06].

Example 3 (Reasoning with equational theories) Consider the signature Σ_H and the equational theory E_H . Show that each of the following equalities hold:

1. $\text{sdec}(k, \text{senc}(k, L)) =_{E_H} L$
2. $\text{pair}(M, N) =_{E_H} \text{pair}(\text{sdec}(k, \text{senc}(k, \text{fst}(\text{pair}(M, N))))), N)$
3. $\text{fst}(\text{sdec}(\text{snd}(\text{pair}(K, L)), \text{senc}(\text{fst}(\text{pair}(L, N)), \text{pair}(M, K)))) =_{E_H} M$

Hint. Figure 1 may be helpful.

Contexts may be used to represent the adversarial environment in which a process is run; that environment provides the data that the process inputs, and consumes the data that it outputs. We define *context* $C[_]$ to be an extended process with a hole. We obtain $C[A]$ as the result of filling $C[_]$'s hole with the extended process A . An *evaluation context* is a context whose hole is not in the scope of a replication, a conditional, an input, or an output. A context $C[_]$ closes A when $C[A]$ is closed.

2.2.1 Structural equivalence

Informally, two processes are structurally equivalent if they model the same thing, but the grammar permits different encodings. For example, to describe a pair of processes A, B running in parallel, the grammar forces us to put one on the left-hand side, and one on the right-hand side of the parallel operator; that is, we have to write either $A \mid B$ or $B \mid A$. These two processes are said to be structurally equivalent. Formally, *structural equivalence* (\equiv) is the smallest equivalence relation on extended processes that is closed by α -conversion of both bound names and bound variables, and closed under application of evaluation contexts such that:

$$\begin{array}{ll}
\text{PAR-0} & A \equiv A \mid 0 \\
\text{PAR-A} & A \mid (B \mid C) \equiv (A \mid B) \mid C \\
\text{PAR-C} & A \mid B \equiv B \mid A \\
\text{REPL} & !P \equiv P \mid !P \\
\\
\text{NEW-0} & \nu n.0 \equiv 0 \\
\text{NEW-C} & \nu u.\nu w.A \equiv \nu w.\nu u.A \\
\text{NEW-PAR} & A \mid \nu u.B \equiv \nu u.(A \mid B) \\
& \text{where } u \notin \text{fv}(A) \cup \text{fn}(A) \\
\\
\text{ALIAS} & \nu x.\{M/x\} \equiv 0 \\
\text{SUBST} & \{M/x\} \mid A \equiv \{M/x\} \mid A\{M/x\} \\
\text{REWRITE} & \{M/x\} \equiv \{N/x\} \\
& \text{where } M =_E N
\end{array}$$

The rules for parallel composition, replication and restriction are self-explanatory. ALIAS enables the introduction of an arbitrary active substitution with restricted scope. SUBST describes the application of an active substitution to a process that it comes into contact with, where $A\{M/x\}$ is A with all free occurrences of x replaced by M (in such a way that no names or variables in M become bound). The final rule, REWRITE, allows terms that are equal modulo the equational theory to be swapped as desired.

Structural equivalence allows every closed extended process A to be rewritten as a substitution and a closed plain process with some restricted names: $A \equiv \nu \tilde{n}.(\{\tilde{M}/\tilde{x}\} \mid P)$ where $\text{fv}(\tilde{M}) = \text{fv}(P) = \emptyset$ and $\tilde{n} \subseteq \text{fn}(\tilde{M})$. It follows immediately that every closed frame φ can be rewritten as a substitution with some restricted names: $\varphi \equiv \nu \tilde{n}.(\{\tilde{M}/\tilde{x}\})$ where $\text{fv}(\tilde{M}) = \emptyset$ and $\tilde{n} \subseteq \text{fn}(\tilde{M})$. We note that the domain of φ is \tilde{x} .

2.2.2 Internal reduction.

A process can be executed without contact with its environment; either because if-statements are evaluated and the then- or else-branch is taken, or because internal subprocesses communicate with each other. The execution of a process with respect to control flow and communication is captured by *internal reduction*. Formally, internal reduction (\rightarrow) is the smallest relation on extended processes closed under structural equivalence and application of evaluation contexts such that:

$$\begin{array}{ll}
\text{COMM} & \bar{c}\langle M \rangle.P \mid c(x).Q \rightarrow P \mid Q\{M/x\} \\
\text{THEN} & \text{if } N = N \text{ then } P \text{ else } Q \rightarrow P \\
\text{ELSE} & \text{if } L = M \text{ then } P \text{ else } Q \rightarrow Q \\
& \text{for ground terms } L, M \text{ where } L \neq_E M
\end{array}$$

We write \rightarrow^* for the reflexive and transitive closure of \rightarrow .

The original presentation [AF01] defined a variant of the communication (COMM) rule on variables rather than terms without loss of generality, since active substitutions were defined over all types and, hence, structural equivalence could be used to allow communication of an arbitrary term instead of a variable. However, as noted in Section 2.1, we restrict active substitutions to terms of base type and to avoid losing generality we define communication (COMM) on terms.

Conditionals (THEN and ELSE) are dependent on the equational theory. The rule THEN is defined for syntactically identical terms N , making it look rather restricted. However, this entails no loss of generality because structural equivalence allows us to replace a term M with N , when $M =_E N$. For instance, given a process $\text{if } M = N \text{ then } P \text{ else } Q$ such that $M =_E N$ and a fresh variable x , we have:

$$\begin{array}{lll}
\text{if } M = N \text{ then } P \text{ else } Q & \equiv & \text{if } M = N \text{ then } P \text{ else } Q \mid 0 & \text{by PAR-0} \\
& \equiv & \text{if } M = N \text{ then } P \text{ else } Q \mid \nu x.\{M/x\} & \text{by ALIAS} \\
& \equiv & \nu x.(\text{if } M = N \text{ then } P \text{ else } Q \mid \{M/x\}) & \text{by NEW-PAR} \\
& \equiv & \nu x.(\{M/x\} \mid \text{if } M = N \text{ then } P \text{ else } Q) & \text{by PAR-C} \\
& \equiv & \nu x.(\{M/x\} \mid \text{if } x = N \text{ then } P \text{ else } Q) & \text{by SUBST} \\
& \equiv & \nu x.(\{N/x\} \mid \text{if } x = N \text{ then } P \text{ else } Q) & \text{by REWRITE} \\
& \equiv & \nu x.(\{N/x\} \mid \text{if } N = N \text{ then } P \text{ else } Q) & \text{by SUBST} \\
& \equiv & \nu x.(\text{if } N = N \text{ then } P \text{ else } Q \mid \{N/x\}) & \text{by PAR-C} \\
& \equiv & \text{if } N = N \text{ then } P \text{ else } Q \mid \nu x.\{N/x\} & \text{by NEW-PAR} \\
& \equiv & \text{if } N = N \text{ then } P \text{ else } Q \mid 0 & \text{by ALIAS} \\
& \equiv & \text{if } N = N \text{ then } P \text{ else } Q & \text{by PAR-0}
\end{array}$$

Structural equivalence may also be required by the ELSE rule to ensure terms L and M are ground by applying active substitutions that appear in the context.

2.2.3 Labelled reductions

The labelled semantics defines a ternary relation written, $A \xrightarrow{\alpha} B$, where α is a *label* of the form $c(M)$, $\bar{c}\langle u \rangle$, or $\nu u.\bar{c}\langle u \rangle$ such that u is either a channel name or a variable of base type.

The transition $A \xrightarrow{c(M)} B$ means that the process A performs an input of the term M from the environment on the channel c , and the resulting process is B . The situation for output is a bit more complicated, since there are several cases. If the item is a free variable x or a free channel name d , then the label $\bar{c}\langle x \rangle$, respectively $\bar{c}\langle d \rangle$, is used. If the item being output is a restricted channel name d , then the label $\nu d.\bar{c}\langle d \rangle$ is used. Finally, if the item is a term M , then the label $\nu x.\bar{c}\langle x \rangle$ is used, after replacing the occurrence of the term M by x and wrapping the process in

$\nu x.(\{M/x\} \mid _)$. The operational semantics of §2.2.1 are extended to include the following rules:

$$\begin{array}{c}
\text{IN} \qquad c(x).P \xrightarrow{c(M)} P\{M/x\} \\
\text{OUT-ATOM} \qquad \bar{c}\langle u \rangle.P \xrightarrow{\bar{c}\langle u \rangle} P \\
\text{OPEN-ATOM} \qquad \frac{A \xrightarrow{\bar{c}\langle u \rangle} A' \quad u \neq c}{\nu u.A \xrightarrow{\nu u.\bar{c}\langle u \rangle} A'} \\
\text{SCOPE} \qquad \frac{A \xrightarrow{\alpha} A' \quad u \text{ does not occur in } \alpha}{\nu u.A \xrightarrow{\alpha} \nu u.A'} \\
\text{PAR} \qquad \frac{A \xrightarrow{\alpha} A' \quad \text{bv}(\alpha) \cap \text{fv}(B) = \text{bn}(\alpha) \cap \text{fn}(B) = \emptyset}{A \mid B \xrightarrow{\alpha} A' \mid B} \\
\text{STRUCT} \qquad \frac{A \equiv B \quad B \xrightarrow{\alpha} B' \quad B' \equiv A'}{A \xrightarrow{\alpha} A'}
\end{array}$$

The rule OPEN-ATOM is used in two ways: 1) to output a restricted channel; and 2) to output a term. The first instance is straight forward and one may prove that $\nu d.\bar{c}\langle d \rangle.P \xrightarrow{\nu d.\bar{c}\langle d \rangle} P$, provided $c \neq d$. The latter is a little more complicated since it requires the use of structural equivalence rules. Observe that $\bar{c}\langle M \rangle.P \xrightarrow{\nu x.\bar{c}\langle x \rangle} P \mid \{M/x\}$, where $x \notin \text{fv}(P)$, by first writing $\bar{c}\langle M \rangle.P$ as $\nu x.(\bar{c}\langle x \rangle.P \mid \{M/x\})$. The full derivation is given below:

$$\frac{\frac{\frac{\frac{\text{OUT-ATOM}}{\bar{c}\langle x \rangle.P \xrightarrow{\bar{c}\langle x \rangle} P}}{\text{PAR}}}{\bar{c}\langle x \rangle.P \mid \{M/x\} \xrightarrow{\bar{c}\langle x \rangle} P \mid \{M/x\}}}{\text{OPEN-ATOM}}}{\bar{c}\langle M \rangle.P \equiv \nu x.(\bar{c}\langle x \rangle.P \mid \{M/x\}) \xrightarrow{\nu x.\bar{c}\langle x \rangle} P \mid \{M/x\} \equiv P \mid \{M/x\}}{\text{STRUCT}} \bar{c}\langle M \rangle.P \xrightarrow{\nu x.\bar{c}\langle x \rangle} P \mid \{M/x\}$$

Note that the fact $x \notin \text{fv}(P)$ is needed for the first of the two structural equivalences in the occurrence of STRUCT.

Example 4 (Labelled semantics) Consider the process $A \triangleq \bar{c}\langle m \rangle.\bar{c}\langle m \rangle$. Show the reductions $A \xrightarrow{\nu x.\bar{c}\langle x \rangle} \xrightarrow{\nu y.\bar{c}\langle y \rangle} \{m/x\} \mid \{m/y\}$. Also show the reductions $A \xrightarrow{\nu x.\bar{c}\langle x \rangle} \xrightarrow{\bar{c}\langle x \rangle} \{m/x\}$. (Hint. The last one involves an application of SUBST.)

2.2.4 Names and the environment

The reduction $\nu \tilde{n}.\bar{c}\langle M \rangle.P \xrightarrow{\nu x.\bar{c}\langle x \rangle} \nu \tilde{n}.(P \mid \{M/x\})$ represents an important idiom of the applied pi calculus. It is the way that an arbitrary term M is output to the environment. This illustrates the way that the environment's use of names is controlled.

If s is a restricted name in a process then the environment cannot use that s to construct a term. Moreover, if the environment uses the name s , then it is not the same as the one in the process. Consider for example the process

$$A \triangleq \nu s.(c(x).\text{if } x = s \text{ then } \bar{c}\langle i_got_s \rangle)$$

This process can never output i_got_s , because no term input as x can be equal to the 'new' s created by the process. More precisely, there is no sequence of reductions

$$A \rightarrow^* \xrightarrow{\alpha} \rightarrow^* \dots \rightarrow^* \xrightarrow{\alpha} \rightarrow^* B \mid \{i_got_s/y\}$$

for some process B and variable y .

Now suppose the process A' creates a new s and outputs some term containing s :

$$A' \triangleq \nu s.(\bar{c}\langle \text{hash}(s) \rangle).B'$$

We have $A' \xrightarrow{\nu x.\bar{c}\langle x \rangle} \nu s.(B' \mid \{\text{hash}(s)/x\})$. This process exposes $\text{hash}(s)$ to the environment, by reference to the handle x . Although the environment still does not have s , it can use $\text{hash}(s)$ in an expression, simply by using x in its place. For example, it can use x to construct input in B' .

Now consider the process A'' which creates a new s , outputs the encryption of s by a (free) key k , and then accepts an input and tests if the input is s .

$$A'' \triangleq \nu s.(\bar{c}\langle \text{senc}(k, s) \rangle).c(x).\text{if } x = s \text{ then } \bar{c}\langle i_got_s \rangle$$

This test can succeed; the process can output i_got_s , as shown by the following execution:

$$\begin{aligned} A'' & \xrightarrow{\nu y.\bar{c}\langle y \rangle} \nu s.(c(x).\text{if } x = s \text{ then } \bar{c}\langle i_got_s \rangle \mid \{\text{senc}(k, s)/y\}) \\ & \xrightarrow{c(\text{sdec}(k, y))} \nu s.(\text{if } \text{sdec}(k, y) = s \text{ then } \bar{c}\langle i_got_s \rangle \mid \{\text{senc}(k, s)/y\}) \\ & \equiv \nu s.(\text{if } \text{sdec}(k, \text{senc}(k, s)) = s \text{ then } \bar{c}\langle i_got_s \rangle \mid \{\text{senc}(k, s)/y\}) \\ & \equiv \nu s.(\text{if } s = s \text{ then } \bar{c}\langle i_got_s \rangle \mid \{\text{senc}(k, s)/y\}) \\ & \rightarrow \nu s.(\bar{c}\langle i_got_s \rangle \mid \{\text{senc}(k, s)/y\}) \\ & \xrightarrow{\nu z.\bar{c}\langle z \rangle} \nu s.(\{\text{senc}(k, s)/y\} \mid \{i_got_s/z\}) \\ & \equiv \nu s.(\{\text{senc}(k, s)/y\}) \mid \{i_got_s/z\} \end{aligned}$$

The first of the equivalences \equiv holds by the rule SUBST; the second one is by REWRITE, using the equation $\text{sdec}(k, \text{senc}(k, s)) = s$. Intuitively, the environment has taken the encrypted output and decrypted it (since the key k is a name known to the environment). It can then input the result to the process, which sees that it is indeed equal to s .

It is instructive to return to the process $A \triangleq \nu s.(c(x).\text{if } x = s \text{ then } \bar{c}\langle i_got_s \rangle)$ above, and consider what happens if the environment gives it the name s as input. As stated earlier, intuitively this s from the environment is considered different from the s that the process constructs. Technically, this is handled as follows. If we try to perform a transition $A \xrightarrow{c(s)} B$, we find that we cannot quite use the expected combination of the rules IN and SCOPE, because IN asks us to perform the substitution

$$\nu s.(\text{if } x = s \text{ then } \bar{c}\langle i_got_s \rangle\{s/x\})$$

which would insert the s from the substitution into the scope of the νs . As previously mentioned, such substitutions are not allowed. We can resolve the situation by α -renaming the bound s first:

$$A \equiv \nu s'.(c(x).\text{if } x = s' \text{ then } \bar{c}\langle i_got_s \rangle).$$

Recall that if $M = N$ then P is an abbreviation of $\text{if } M = N \text{ then } P \text{ else } 0$ and hence we have the reductions:

$$\nu s'.(c(x).\text{if } x = s' \text{ then } \bar{c}\langle i_got_s \rangle) \xrightarrow{c(s)} \text{if } s = s' \text{ then } \bar{c}\langle i_got_s \rangle \rightarrow 0.$$

3 Secrecy and correspondence properties

This section presents formal definitions of secrecy and correspondence in the presence of an adversary who has full control of the network. The attacker can therefore eavesdrop, replay, inject and block messages. Formally the attacker is captured as an arbitrary process and is sometimes called the Dolev-Yao adversary [DY83].

3.1 Secrecy

Intuitively, a protocol preserves the secrecy of some term M if an adversary cannot obtain M by constructing it from the outputs of the protocol [AB05a]. We formalise the adversary as a process running in parallel with the protocol, that after constructing M outputs it on a public channel. The adversary process does not have any of the protocol's secrets.

A term may refer to names inside a process; to analyse secrecy of the term, it is important that these names appear unambiguously in the process. Otherwise it is not clear which name the term refers to. For example, if a name in the term is both bound and free in the process, then it is not clear whether the term refers to the free instance, or the bound instance. We formalise that by requiring the process to be “name distinct” for the names mentioned in the term.

Definition 1 (name-distinct for \tilde{m}) *A closed plain process P is name-distinct for a set of names \tilde{m} , if $\tilde{m} \cap \text{fn}(P) \cap \text{bn}(P) = \emptyset$ and for each name $n \in \tilde{m} \cap \text{bn}(P)$ there is exactly one name restriction νn occurring in P , and the restriction is not under replication “!*”.

Definition 2 (Can output) *A plain process P can output the term M if there exists an evaluation context $C[_]$, a channel name $c \notin \text{bn}(C)$ and process R such that the reduction $P \rightarrow^* C[\bar{c}\langle M \rangle.R]$ holds with no alpha-renaming of the names in $\text{fn}(M)$.*

In Definition 2, the process $C[\bar{c}\langle M \rangle.R]$ is capable of outputting the term M on the free channel c in one further step. Note that the definition forbids renaming names in M during the reduction $P \rightarrow^* C[\bar{c}\langle M \rangle.R]$, because we must not change which names in P are referred to by M . Thus, the process $\nu b.\bar{c}\langle b \rangle$ cannot output s ; if we allowed renaming, we could write the process as $\nu s.\bar{c}\langle s \rangle$ and it would be able to output s .

Definition 3 (Reachability-based secrecy) *Let M be a term, and P be a closed plain process that is name-distinct for $\text{fn}(M)$. Then P preserves the reachability-based secrecy of M if there is no plain process I such that $(\text{fn}(I) \cup \text{bn}(I)) \cap \text{bn}(P) = \emptyset$ and $P \mid I$ can output M .*

In the definition above, I is the adversary (or *intruder*) process. Typically, it is built in order to receive the outputs from P , and then possibly to construct from them the secret term M , and output it. If there is no such intruder, then P keeps M secret. This definition is based on the one in [AB05a], but extends it to cope with bound names.

Example 5 (Reasoning with secrecy) *Consider the process P corresponding to the Handshake protocol (Section 2.1.1) and the equational theory E_H . We show that P does not preserve the secrecy of s , by presenting an adversarial process*

$$\begin{aligned} I \triangleq & c(y_pk).\bar{c}\langle \text{pk}(sk_M) \rangle.c(x). \\ & \bar{c}\langle \text{aenc}(y_pk, \text{adec}(sk_M, x)) \rangle.c(z). \\ & \bar{c}\langle \text{snd}(\text{sdec}(\text{getmsg}(\text{adec}(sk_M, x)), z)) \rangle \end{aligned}$$

and demonstrating that $P \mid I$ can evolve to a process that can output s on a public channel. To aid readability, we apply all the substitutions denoted by ‘let’ occurring in P , and we use the context

$$C[_] \triangleq \nu sk_S.\nu sk_C.\nu s.(- \mid !P_S \mid !P_C)$$

and write $P \equiv C[\bar{c}\langle \text{pk}(sk_S) \rangle \mid \bar{c}\langle \text{pk}(sk_C) \rangle \mid P_S \mid P_C]$. Since $sk_S, sk_C, s \notin \text{fn}(I) \cup \text{fv}(I)$, we have

$$\begin{aligned} P \mid I \equiv & C[\\ & \bar{c}\langle \text{pk}(sk_S) \rangle \mid \bar{c}\langle \text{pk}(sk_C) \rangle \\ & \mid c(x_pk).\nu k.\bar{c}\langle \text{aenc}(x_pk, \text{sign}(sk_S, k)) \rangle. \\ & \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\ & \mid c(y).\text{if checksign}(\text{pk}(sk_S), \text{adec}(sk_C, y)) = \text{true} \text{ then} \\ & \quad \bar{c}\langle \text{senc}(\text{getmsg}(\text{adec}(sk_C, y)), \text{pair}(\text{tag}, s)) \rangle \\ & \mid I \\ &] \end{aligned}$$

Figure 2 Execution of Handshake protocol (internal reduction)

$$\begin{array}{l}
1. \quad P \mid I \rightarrow C[\bar{c}(\text{pk}(sk_S)) \\
\quad \quad | \quad c(x.pk).\nu k.\bar{c}(\text{aenc}(x.pk, \text{sign}(sk_S, k))). \\
\quad \quad | \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\
\quad \quad | \quad c(y).\text{if checksign}(\text{pk}(sk_S), \text{adec}(sk_C, y)) = \text{true} \text{ then} \\
\quad \quad | \quad \quad \bar{c}(\text{senc}(\text{getmsg}(\text{adec}(sk_C, y)), \text{pair}(\text{tag}, s))) \\
\quad \quad | \quad \bar{c}(\text{pk}(sk_M)).c(x).\bar{c}(\text{aenc}(\text{pk}(sk_C), \text{adec}(sk_M, x))).c(z). \\
\quad \quad | \quad \bar{c}(\text{snd}(\text{sdec}(\text{getmsg}(\text{adec}(sk_M, x)), z)))] \\
2. \quad \rightarrow C[\bar{c}(\text{pk}(sk_S)) \\
\quad \quad | \quad \nu k.\bar{c}(\text{aenc}(\text{pk}(sk_M), \text{sign}(sk_S, k))). \\
\quad \quad | \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\
\quad \quad | \quad c(y).\text{if checksign}(\text{pk}(sk_S), \text{adec}(sk_C, y)) = \text{true} \text{ then} \\
\quad \quad | \quad \quad \bar{c}(\text{senc}(\text{getmsg}(\text{adec}(sk_C, y)), \text{pair}(\text{tag}, s))) \\
\quad \quad | \quad c(x).\bar{c}(\text{aenc}(\text{pk}(sk_C), \text{adec}(sk_M, x))).c(z). \\
\quad \quad | \quad \bar{c}(\text{snd}(\text{sdec}(\text{getmsg}(\text{adec}(sk_M, x)), z)))] \\
3. \quad \rightarrow \nu k.C[\bar{c}(\text{pk}(sk_S)) \\
\quad \quad | \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\
\quad \quad | \quad c(y).\text{if checksign}(\text{pk}(sk_S), \text{adec}(sk_C, y)) = \text{true} \text{ then} \\
\quad \quad | \quad \quad \bar{c}(\text{senc}(\text{getmsg}(\text{adec}(sk_C, y)), \text{pair}(\text{tag}, s))) \\
\quad \quad | \quad \bar{c}(\text{aenc}(\text{pk}(sk_C), \text{adec}(sk_M, \text{aenc}(\text{pk}(sk_M), \text{sign}(sk_S, k))))).c(z). \\
\quad \quad | \quad \bar{c}(\text{snd}(\text{sdec}(\text{getmsg}(\text{adec}(sk_M, \text{aenc}(\text{pk}(sk_M), \text{sign}(sk_S, k))))), z)))] \\
4. \quad \rightarrow \nu k.C[\bar{c}(\text{pk}(sk_S)) \\
\quad \quad | \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\
\quad \quad | \quad \text{if checksign}(\text{pk}(sk_S), \text{adec}(sk_C, \text{aenc}(\text{pk}(sk_C), \\
\quad \quad \quad \text{adec}(sk_M, \text{aenc}(\text{pk}(sk_M), \text{sign}(sk_S, k)))))) = \text{true} \text{ then} \\
\quad \quad | \quad \quad \bar{c}(\text{senc}(\text{getmsg}(\text{adec}(sk_C, \text{aenc}(\text{pk}(sk_C), \\
\quad \quad \quad \text{adec}(sk_M, \text{aenc}(\text{pk}(sk_M), \text{sign}(sk_S, k)))))), \text{pair}(\text{tag}, s))) \\
\quad \quad | \quad c(z). \\
\quad \quad | \quad \bar{c}(\text{snd}(\text{sdec}(\text{getmsg}(\text{adec}(sk_M, \text{aenc}(\text{pk}(sk_M), \text{sign}(sk_S, k))))), z)))] \\
5. \quad \equiv \nu k.C[\bar{c}(\text{pk}(sk_S)) \\
\quad \quad | \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\
\quad \quad | \quad \text{if true} = \text{true} \text{ then} \\
\quad \quad | \quad \quad \bar{c}(\text{senc}(k, \text{pair}(\text{tag}, s))) \\
\quad \quad | \quad c(z).\bar{c}(\text{snd}(\text{sdec}(k, z)))] \\
6. \quad \rightarrow\rightarrow \nu k.C[\bar{c}(\text{pk}(sk_S)) \\
\quad \quad | \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\
\quad \quad | \quad \bar{c}(\text{snd}(\text{sdec}(k, \text{senc}(k, \text{pair}(\text{tag}, s)))))] \\
7. \quad \equiv \nu k.C[\bar{c}(\text{pk}(sk_S)) \\
\quad \quad | \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\
\quad \quad | \quad \bar{c}(s)]
\end{array}$$

The execution path which results in the adversary being able to output s on a public channel is shown in Figure 2. Intuitively, the following sequence takes place (the item numbers correspond to the transitions in the figure):

1. C 's public key $\text{pk}(sk_C)$ is published and hence made available to the adversary, who inputs it as y_{pk} .
2. The adversary provides her public key $\text{pk}(sk_M)$ as S 's interlocutor.
3. S outputs her signature on k , encrypted for the adversary, and the adversary inputs it as x . (Observe that closure under structural equivalence is used to move the name restriction νk outside of the context; more precisely, the rules NEW-C, NEW-PAR are applied.)

4. The adversary gives to \mathcal{C} the value k signed by \mathcal{S} , this time encrypted for \mathcal{C} .
5. The process is further rewritten using structural equivalence (essentially several occurrences of the rule REWRITE with respect to E_H). (Note that this rewriting could have been included in the previous step because internal reduction is closed under structural equivalence, but we explicitly present it for clarity.)
6. The conditional is evaluated and \mathcal{C} sends the value $\text{senc}(k, \text{pair}(\text{tag}, s))$ to the adversary.
7. Observe that $\bar{c}(\text{snd}(\text{sdec}(k, \text{senc}(k, \text{pair}(\text{tag}, s)))))) \equiv \bar{c}(s)$ and hence the adversary can obtain s and publish the value on a public channel.

3.2 Correspondence properties

Correspondence properties are used to capture relationships between events that can be expressed in the form “if an event e has been executed then event e' has been previously executed.” Moreover, these events may contain arguments, which allow relationships between the arguments of events to be expressed. To reason with correspondence properties we annotate processes with *events*, marking important stages reached by the protocol which do not otherwise affect behaviour. Events are analogous to breakpoints used in software development. In this chapter, we only consider basic correspondence properties, which are sufficient to model authentication. More elaborate formalisms can be found in [ABF07, Bla09]. Events are message outputs $\bar{f}\langle M \rangle$ where f is an *event channel* (a name in a particular set, disjoint from the set of ordinary channel names). In labelled transitions, output labels for events use *event variables* e . Those event variables are not allowed to appear in input labels $u(M)$, so the adversary cannot use them. (This condition is important, since events are added just for checking correspondence properties; in particular an event $\bar{f}\langle M \rangle$ does not reveal M to the adversary.) Hence, the execution of the process P after inserting events $\bar{f}\langle M \rangle$ is the execution of P without events, plus the recording of events using labels $\nu e.\bar{f}\langle e \rangle$ and active substitutions $\{M/e\}$.

Definition 4 (Correspondence property) A correspondence property is a formula of the form: $\bar{f}\langle M \rangle \rightsquigarrow \bar{g}\langle N \rangle$.

The property asserts that if the event f has been executed in a trace with parameters M , then the event g must have been previously executed with parameters N .

Authentication can be captured as a correspondence property. Recall that in addition to the secrecy property mentioned for the Handshake protocol in Section 1, there were also authentication properties. The protocol is intended to ensure that if \mathcal{S} thinks she executes the protocol with \mathcal{C} , then she really does so, and vice versa. When we say ‘she thinks’ that she executes it with \mathcal{C} , we mean that the data she receives indicates that fact. Accordingly we annotate the Handshake protocol with events. To capture equality tests within events we include the binary function eq in the signature Σ_H and extend the equational theory E_H with the equation $\text{eq}(x, x) = \text{true}$.

Example 6 (Annotated Handshake protocol) The annotated Handshake protocol is presented below

$$\begin{aligned}
P &\triangleq \nu sk_S. \nu sk_C. \nu s. \\
&\quad \text{let } pk_S = \text{pk}(sk_S) \text{ in let } pk_C = \text{pk}(sk_C) \text{ in} \\
&\quad (\bar{c}\langle pk_S \rangle \mid \bar{c}\langle pk_C \rangle \mid !P_S \mid !P_C) \\
P_S &\triangleq c(x.pk). \nu k. \overline{\text{startedS}}\langle \text{pair}(x.pk, k) \rangle \\
&\quad \bar{c}\langle \text{aenc}(x.pk, \text{sign}(sk_S, k)) \rangle. \\
&\quad c(z). \text{if } \text{fst}(\text{sdec}(k, z)) = \text{tag} \text{ then} \\
&\quad \overline{\text{completedS}}\langle \text{pair}(k, \text{eq}(x.pk, pk_C)) \rangle. Q
\end{aligned}$$

$$P_C \triangleq \frac{c(y).\text{let } y' = \text{adec}(sk_C, y) \text{ in let } y_k = \text{getmsg}(y') \text{ in}}{\text{started}\overline{C}\langle y_k \rangle} \text{if } \text{checksign}(pk_S, y') = \text{true} \text{ then} \\ \frac{\overline{c}\langle \text{senc}(y_k, \text{pair}(tag, s)) \rangle}{\text{completed}\overline{C}\langle \text{pair}(pk_C, y_k) \rangle}$$

where the four events are interpreted as follows:

- $\overline{\text{started}}\overline{S}\langle \text{pair}(x_pk, k) \rangle$ means that \mathcal{S} considers she has started the protocol with an interlocutor whose public key is x_pk , and she has proposed k as the session key.
- $\overline{\text{started}}\overline{C}\langle y_k \rangle$ means that \mathcal{C} considers he has started the protocol with the session key y_k .
- $\overline{\text{completed}}\overline{S}\langle \text{pair}(k, t) \rangle$ means that \mathcal{S} believes she has completed the protocol with session key k and if $t =_{E_H} \text{true}$, then the protocol was completed with \mathcal{C} .
- $\overline{\text{completed}}\overline{C}\langle \text{pair}(pk_C, y_k) \rangle$ means that \mathcal{C} considers he has successfully completed the protocol with \mathcal{S} using session key y_k , where pk_C is his public key.

Correspondence properties can now be defined to allow the analysis of authentication. Recall that the client \mathcal{C} is only willing to share her secret with the server \mathcal{S} . We formalise the authentication of \mathcal{C} using the correspondence property

$$\overline{\text{completed}}\overline{C}\langle \text{pair}(x, y) \rangle \rightsquigarrow \overline{\text{started}}\overline{S}\langle \text{pair}(x, y) \rangle.$$

In comparison, \mathcal{S} is willing to run the protocol with any other principal. The correspondence property says that if she believes \mathcal{C} was her interlocutor, then \mathcal{C} must have completed the protocol with the suggested key. This is formalised as:

$$\overline{\text{completed}}\overline{S}\langle \text{pair}(y, \text{true}) \rangle \rightsquigarrow \overline{\text{started}}\overline{C}\langle y \rangle$$

The subtle differences between the two correspondence properties is due to the differing authentication properties expected by participants \mathcal{S} and \mathcal{C} .

Formally we define the validity of a correspondence property in Definition 5. Intuitively, it ensures that if the event f is executed, then the event g must have been previously executed. Moreover, the parametrisation of the events must satisfy any relationships defined by M, N ; that is, the variables $\text{fv}(M) \cap \text{fv}(N)$ must be parametrised in the same way.

Definition 5 (Validity of correspondence property) Let E be an equational theory, and A_0 an extended process. We say that A_0 satisfies the correspondence property $\overline{f}\langle M \rangle \rightsquigarrow \overline{g}\langle N \rangle$ if for all execution paths

$$A_0 \rightarrow^* \xrightarrow{\alpha_1} \rightarrow^* A_1 \rightarrow^* \xrightarrow{\alpha_2} \rightarrow^* \dots \rightarrow^* \xrightarrow{\alpha_n} \rightarrow^* A_n,$$

and all index $i \in \mathbb{N}$, substitution σ and variable e such that $\alpha_i = \nu e.\overline{f}\langle e \rangle$ and $e\varphi(A_i) =_E M\sigma$, there exists $j \in \mathbb{N}$ and e' such that $\alpha_j = \nu e'.\overline{g}\langle e' \rangle$, $e'\varphi(A_j) =_E N\sigma$ and $j < i$.

Example 7 (Reasoning with correspondence) Consider the annotated Handshake protocol (Example 6). The first correspondence property, namely

$$\overline{\text{completed}}\overline{C}\langle \text{pair}(x, y) \rangle \rightsquigarrow \overline{\text{started}}\overline{S}\langle \text{pair}(x, y) \rangle$$

is not valid. This can be observed by constructing an execution path in which the events $\overline{\text{completed}}\overline{C}$ and $\overline{\text{started}}\overline{S}$ both occur, but with different arguments. To aid readability, we apply some of the substitutions denoted by ‘let’ occurring in P , and we reuse the context $C[_] \triangleq \nu sk_S.\nu sk_C.\nu s.(_ | !P_S | !P_C)$ from Example 5. The execution path is shown in Figures 3 & 4. Intuitively, the following sequence of actions takes place:

1. \mathcal{C} 's public key $pk(sk_C)$ is output, using the handle y_pk . This public key is now available to the environment.

Figure 3 Execution of Handshake protocol, part I (labelled semantics)

$$\begin{array}{l}
1. \quad P \xrightarrow{\nu y_pk.\bar{c}(y_pk)} C[\bar{c}\langle\text{pk}(sk_S)\rangle \mid \{\text{pk}(sk_C)/y_pk\} \\
\quad \quad \quad \mid c(x_pk).\nu k.\overline{\text{startedS}}\langle\text{pair}(x_pk, k)\rangle. \\
\quad \quad \quad \quad \bar{c}\langle\text{aenc}(x_pk, \text{sign}(sk_S, k))\rangle. \\
\quad \quad \quad \quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then} \\
\quad \quad \quad \quad \overline{\text{completedS}}\langle\text{pair}(k, \text{eq}(x_pk, \text{pk}(sk_C)))\rangle.Q \\
\quad \quad \quad \mid c(y). \\
\quad \quad \quad \quad \text{let } y' = \text{adec}(sk_C, y) \text{ in} \\
\quad \quad \quad \quad \text{let } y_k = \text{getmsg}(y') \text{ in} \\
\quad \quad \quad \quad \overline{\text{startedC}}\langle y_k \rangle. \\
\quad \quad \quad \quad \text{if checksign}(\text{pk}(sk_S), y') = \text{true} \text{ then} \\
\quad \quad \quad \quad \bar{c}\langle\text{senc}(y_k, \text{pair}(\text{tag}, s))\rangle \\
\quad \quad \quad \quad \overline{\text{completedC}}\langle\text{pair}(\text{pk}(sk_C), y_k)\rangle] \\
2. \quad \quad \quad \xrightarrow{c(\text{pk}(sk_M))} \\
3. \quad \quad \quad \xrightarrow{\nu e_1.\overline{\text{startedS}}\langle e_1 \rangle} \\
4. \quad \quad \quad \xrightarrow{\nu x.\bar{c}(x)} \nu k.C[\bar{c}\langle\text{pk}(sk_S)\rangle \mid \{\text{pk}(sk_C)/y_pk\} \\
\quad \quad \quad \mid \{\text{pair}(\text{pk}(sk_M), k)/e_1\} \\
\quad \quad \quad \mid \{\text{aenc}(\text{pk}(sk_M), \text{sign}(sk_S, k))/x\} \\
\quad \quad \quad \mid c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then} \\
\quad \quad \quad \quad \overline{\text{completedS}}\langle\text{pair}(k, \text{eq}(\text{pk}(sk_M), \text{pk}(sk_C)))\rangle.Q \\
\quad \quad \quad \mid c(y). \\
\quad \quad \quad \quad \text{let } y' = \text{adec}(sk_C, y) \text{ in} \\
\quad \quad \quad \quad \text{let } y_k = \text{getmsg}(y') \text{ in} \\
\quad \quad \quad \quad \overline{\text{startedC}}\langle y_k \rangle. \\
\quad \quad \quad \quad \text{if checksign}(\text{pk}(sk_S), y') = \text{true} \text{ then} \\
\quad \quad \quad \quad \bar{c}\langle\text{senc}(y_k, \text{pair}(\text{tag}, s))\rangle \\
\quad \quad \quad \quad \overline{\text{completedC}}\langle\text{pair}(\text{pk}(sk_C), y_k)\rangle]
\end{array}$$

2. The environment provides the public key $\text{pk}(sk_M)$ as S 's interlocutor.
3. The event $\overline{\text{startedS}}\langle\text{pair}(x_pk, k)\rangle$ is executed with the environment's public key $\text{pk}(sk_M)$ assigned to parameter x_pk .
4. S outputs as x the value k signed by S and encrypted for the environment.
5. The environment gives to C the value k signed by S , this time encrypted for C .
6. The process is rewritten using structural equivalence (essentially several occurrences of the rules SUBST and REWRITE).
7. The event $\overline{\text{startedC}}\langle y_k \rangle$ is executed with k as parameter y_k .
8. The conditional is trivially evaluated and C outputs the value $\text{senc}(k, \text{pair}(\text{tag}, s))$ as z .
9. Finally, the event $\overline{\text{completedC}}\langle\text{pair}(\text{pk}(sk_C), y_k)\rangle$ is executed with respect to the value k assigned to parameter y_k .

Thus, the conditions of Definition 5 are violated; we have $\overline{\text{completedC}}\langle e_3 \rangle$, and although we do have a previous $\overline{\text{startedS}}\langle e_1 \rangle$, it does not have the same arguments as can be observed from the frame.

Figure 4 Execution of Handshake protocol, part II (labelled semantics)

$$\begin{array}{l}
5. \quad \frac{c(\overline{\text{aenc}}(y_pk, \overline{\text{adec}}(sk_M, x)))}{\nu k. C[\overline{c}(\text{pk}(sk_S)) \mid \{\text{pk}(sk_C)/y_pk\} \\
\quad \mid \{\text{pair}(\text{pk}(sk_M), k)/e_1\} \\
\quad \mid \{\overline{\text{aenc}}(\text{pk}(sk_M), \text{sign}(sk_S, k))/x\} \\
\quad \mid c(z).\text{if fst}(\overline{\text{sdec}}(k, z)) = \text{tag then} \\
\quad \quad \overline{\text{completed}}S\langle \text{pair}(k, \text{eq}(\text{pk}(sk_M), \text{pk}(sk_C))) \rangle.Q \\
\quad \mid \text{let } y' = \overline{\text{adec}}(sk_C, \overline{\text{aenc}}(y_pk, \overline{\text{adec}}(sk_M, x))) \text{ in} \\
\quad \mid \text{let } y_k = \overline{\text{getmsg}}(y') \text{ in} \\
\quad \quad \overline{\text{started}}C\langle y_k \rangle. \\
\quad \quad \text{if } \overline{\text{checksign}}(\text{pk}(sk_S), y') = \text{true then} \\
\quad \quad \quad \overline{c}\langle \overline{\text{senc}}(y_k, \text{pair}(\text{tag}, s)) \rangle \\
\quad \quad \quad \overline{\text{completed}}C\langle \text{pair}(\text{pk}(sk_C), y_k) \rangle] \\
6. \quad \equiv \nu k. C[\overline{c}(\text{pk}(sk_S)) \mid \{\text{pk}(sk_C)/y_pk\} \\
\quad \mid \{\text{pair}(\text{pk}(sk_M), k)/e_1\} \\
\quad \mid \{\overline{\text{aenc}}(\text{pk}(sk_M), \text{sign}(sk_S, k))/x\} \\
\quad \mid c(z).\text{if fst}(\overline{\text{sdec}}(k, z)) = \text{tag then} \\
\quad \quad \overline{\text{completed}}S\langle \text{pair}(k, \text{eq}(\text{pk}(sk_M), \text{pk}(sk_C))) \rangle.Q \\
\quad \mid \overline{\text{started}}C\langle k \rangle. \\
\quad \mid \text{if } \text{true} = \text{true then} \\
\quad \quad \overline{c}\langle \overline{\text{senc}}(k, \text{pair}(\text{tag}, s)) \rangle \\
\quad \quad \overline{\text{completed}}C\langle \text{pair}(\text{pk}(sk_C), k) \rangle] \\
7. \quad \frac{\nu e_2. \overline{\text{started}}C\langle e_2 \rangle}{\nu z. \overline{c}\langle z \rangle} \\
8. \quad \rightarrow \nu z. \overline{c}\langle z \rangle \\
9. \quad \frac{\nu e_3. \overline{\text{completed}}C\langle e_3 \rangle}{\nu k. C[\overline{c}(\text{pk}(sk_S)) \mid \{\text{pk}(sk_C)/y_pk\} \\
\quad \mid \{\text{pair}(\text{pk}(sk_M), k)/e_1\} \\
\quad \mid \{\overline{\text{aenc}}(\text{pk}(sk_M), \text{sign}(sk_S, k))/x\} \\
\quad \mid c(z).\text{if fst}(\overline{\text{sdec}}(k, z)) = \text{tag then} \\
\quad \quad \overline{\text{completed}}S\langle \text{pair}(k, \text{eq}(\text{pk}(sk_M), \text{pk}(sk_C))) \rangle.Q \\
\quad \mid \{k/e_2\} \\
\quad \mid \{\overline{\text{senc}}(k, \text{pair}(\text{tag}, s))/z\} \\
\quad \mid \{\text{pair}(\text{pk}(sk_C), k)/e_3\}]}
\end{array}$$

3.2.1 Injective correspondence properties

The definition of correspondence we have just discussed is insufficient to capture injective relationships between events; making it unsuitable for certain authentication properties. For example, consider a financial transaction in which a server requests payment from a client; the server should only complete a transaction, when that transaction was started by the client. (If this were not the case, the client could be charged for several transactions, even if the client only started one.) The situation is similar for access control and other scenarios. Further discussion can be found in [Low97].

Example 8 Consider the process

$$\overline{\text{start}}\langle n \rangle. \overline{\text{complete}}\langle n \rangle. \overline{\text{complete}}\langle n \rangle$$

which satisfies the correspondence $\overline{\text{complete}}\langle x \rangle \rightsquigarrow \overline{\text{start}}\langle x \rangle$, but permits two occurrences of the event $\overline{\text{complete}}\langle n \rangle$ to be matched by a single event $\overline{\text{start}}\langle n \rangle$.

Injective correspondence properties are denoted by the presence of the keyword inj.

Definition 6 (Injective correspondence property) An injective correspondence property is a formula of the form: $\overline{f}\langle M \rangle \rightsquigarrow \text{inj } \overline{g}\langle N \rangle$.

Informally, it asserts that if a process executes event f , then there is a distinct earlier occurrence of the event g and any relationship between the event parameters is satisfied. It follows immediately that the number of occurrences of the label $\nu e.\bar{g}\langle e \rangle$ is greater than, or equal to, the number of occurrences of $\nu e'.\bar{f}\langle e' \rangle$ for some event variables e, e' .

Definition 7 (Validity of injective correspondence property) *Let E be an equational theory, and A_0 an extended process. We say that A_0 satisfies the injective correspondence property $\bar{f}\langle M \rangle \rightsquigarrow \text{inj } \bar{g}\langle N \rangle$ if for all execution paths*

$$A_0 \rightarrow^* \xrightarrow{\alpha_1} \rightarrow^* A_1 \rightarrow^* \xrightarrow{\alpha_2} \rightarrow^* \dots \rightarrow^* \xrightarrow{\alpha_n} \rightarrow^* A_n,$$

there exists a partial injective function $h : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ such that for all $i \in \{1, \dots, n\}$, substitution σ and variable e such that $\alpha_i = \nu e.\bar{f}\langle e \rangle$ and $e\varphi(A_i) =_E M\sigma$, then the following conditions are satisfied: (1) $h(i)$ is defined; (2) $\alpha_{h(i)} = \nu e'.\bar{g}\langle e' \rangle$ for some e' such that $e'\varphi(A_{h(i)}) =_E N\sigma$; and, (3) $h(i) < i$.

Returning to our Handshake protocol, observe that authentication of \mathcal{C} is injective; that is, if \mathcal{S} reaches the end of the protocol with the belief that she has session key k with interlocutor \mathcal{C} , then she has indeed done so; and moreover, \mathcal{C} has done so in this session (that is, there is an injective relationship).

Example 9 (Reasoning with correspondence, revisited) *Consider the annotated Handshake protocol and observe the injective correspondence property*

$$\overline{\text{completedS}}\langle \text{pair}(y, \text{true}) \rangle \rightsquigarrow \text{inj } \overline{\text{startedC}}\langle y \rangle$$

is valid. To show this conclusively, one has to consider all the possible execution paths, and show that for each of them, whenever there is a label $\nu e.\overline{\text{completedS}}\langle e \rangle$ associated with the substitution $\{\text{pair}(M, \text{true})/e\}$, then there is an earlier label $\nu e'.\overline{\text{startedC}}\langle e' \rangle$ and substitution $\{M/e'\}$; such that no two occurrences of the event $\overline{\text{completedS}}$ are mapped to the same event $\overline{\text{startedC}}$. ProVerif can be used to prove this result.

3.3 Fixing the Handshake protocol

As mentioned in Section 1, the man-in-the-middle attack that is illustrated in Figure 2 can be avoided by putting the participants' public keys along with k in the signature formed by \mathcal{S} . The resulting process looks like this:

$$\begin{aligned} P &\triangleq \nu sk_S.\nu sk_C.\nu s. \\ &\quad \text{let } pk_S = \text{pk}(sk_S) \text{ in let } pk_C = \text{pk}(sk_C) \text{ in} \\ &\quad (\bar{c}\langle pk_S \rangle \mid \bar{c}\langle pk_C \rangle \mid !P_S \mid !P_C) \\ P_S &\triangleq c(x_pk).\nu k.\bar{c}\langle \text{aenc}(x_pk, \text{sign}(sk_S, \text{pair}(pk_S, \text{pair}(x_pk, k)))) \rangle). \\ &\quad c(z).\text{if fst}(\text{sdec}(k, z)) = \text{tag} \text{ then } Q \\ P_C &\triangleq c(y).\text{let } y' = \text{adec}(sk_C, y) \text{ in let } y'' = \text{snd}(\text{getmsg}(y')) \text{ in} \\ &\quad \text{if fst}(y'') = pk_C \text{ then} \\ &\quad \text{if checksign}(pk_S, y') = \text{true} \text{ then} \\ &\quad \bar{c}\langle \text{senc}(\text{snd}(y''), \text{pair}(\text{tag}, s)) \rangle \end{aligned}$$

As an exercise the reader may like to identify exactly which step of the attack of Figure 2 fails for the revised protocol.

We remark that the injective correspondence property $\overline{\text{completedC}}\langle \text{pair}(x, y) \rangle \rightsquigarrow \text{inj } \overline{\text{startedS}}\langle \text{pair}(x, y) \rangle$ does not hold, because the client cannot ascertain as to whether the initial message he receives is fresh. However, this does not affect the security of the protocol's high-level objective which is to enable the client \mathcal{C} to share the secret s with the server \mathcal{S} .

4 Equivalence properties

The notion of indistinguishability is a powerful concept which allows us to reason about complex properties that cannot be expressed as secrecy or correspondence properties. Intuitively, two processes are said to be equivalent if an observer has no way to tell them apart. The processes may be handling different data, and internally performing quite different computations, but they look the same to an external observer. This notion allows us to define strong notions of secrecy and privacy.

A natural starting point for defining observational equivalence says that processes A and B are equivalent if they can output on the same channels, no matter what the context they are placed inside. Formally we write $A \Downarrow c$ when A can evolve to a process that can send a message on channel c , that is, when $A \rightarrow^* C[\bar{c}\langle M \rangle.P]$ for some term M , process P and evaluation context $C[_]$ that does not bind c . This then allows us to characterise equivalence as follows: for all contexts $C[_]$, we have $C[A] \Downarrow c$ if and only if $C[B] \Downarrow c$. One might think that this condition is too weak – it should say that $C[A]$ and $C[B]$ output the *same term* on the channel c – but that is not necessary, since if two processes output different terms then a context that inspects the output could be constructed to distinguish them.

4.1 Observational equivalence

In practice, the definition motivated above is hard to compute, and instead a stronger definition is often used. The stronger definition, called observational equivalence, has a recursive character. Roughly speaking, processes A and B are said to be observationally equivalent if they can output on the same channel for all contexts they are placed inside (as before); and also, for all contexts and every step made by A inside the context, there exists a step that B can make (inside the context), such that the resulting pair of processes are observationally equivalent (and vice-versa). We avoid the circularity of this informal characterisation, and define it rigorously as follows.

Definition 8 (Observational equivalence) *Observational equivalence (\approx) is the largest symmetric relation \mathcal{R} between closed extended processes with the same domain such that $A \mathcal{R} B$ implies:*

1. if $A \Downarrow c$, then $B \Downarrow c$;
2. if $A \rightarrow^* A'$, then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' ;
3. $C[A] \mathcal{R} C[B]$ for all closing evaluation contexts $C[_]$.

Note that we insist that \mathcal{R} is symmetric, so we can state the Conditions 1 and 2 in a non-symmetric way (the symmetric part is guaranteed by the symmetry of \mathcal{R}). Note also the added condition that the two processes have the same domain; this captures the ‘observable’ difference between processes in which the value of a variable appearing in one frame, is not defined in the other.

A classical example illustrates the difference between observational equivalence and the notion of equivalence mentioned at the start of Section 4 (sometimes called *trace equivalence*). Intuitively, trace equivalence checks that the outputs on the traces allowed by $C[A]$ are the same as those of $C[B]$, but it does not enforce that *decisions* inside A and B occur at the same point. Let us use the notation $A_1 + A_2$ to mean the non-deterministic choice of A_1 or A_2 . This is expressible in applied pi, as follows:

$$A_1 + A_2 \triangleq \nu a.(\bar{a}\langle \text{left} \rangle \mid \bar{a}\langle \text{right} \rangle \mid a(x).\text{if } x = \text{left} \text{ then } A_1 \text{ else } A_2)$$

Here, *left* and *right* are names, and a scheduler chooses which message output $\bar{a}\langle \text{left} \rangle$ or $\bar{a}\langle \text{right} \rangle$ to perform, and hence whether A_1 or A_2 runs. Note that only one of A_1 and A_2 runs, in contrast with $A_1 \mid A_2$; moreover, only one of the outputs $\bar{a}\langle \text{left} \rangle$, $\bar{a}\langle \text{right} \rangle$ may occur, because there is only one input $a(x)$, and a is restricted. Now consider the processes $A \triangleq \bar{d}\langle b_1 \rangle.\bar{d}\langle b_2 \rangle + \bar{d}\langle b_1 \rangle.\bar{d}\langle b_3 \rangle$ and $B \triangleq \bar{d}\langle b_1 \rangle.(\bar{d}\langle b_2 \rangle + \bar{d}\langle b_3 \rangle)$. One can see that they are trace equivalent, by experimenting with a

variety of contexts. However, in A the decision between outputting b_2 or b_3 is taken earlier than it is in B . Observational equivalence captures this as an observable difference. Consider the context $C[-] = _ | d(y)$. Although $C[A]$ and $C[B]$ have the same output capabilities, $C[B]$ can evolve to a process (namely $\bar{d}(b_2) + \bar{d}(b_3)$) which is not equivalent to any process that $C[A]$ can evolve to (namely, the processes $\bar{d}(b_2)$ or $\bar{d}(b_3)$).

4.2 Labelled bisimilarity

The quantification over contexts makes the definition of observational equivalence hard to use in practice. Therefore, labelled bisimilarity is introduced, which is more suitable for both manual and automatic reasoning. Labelled bisimilarity relies on an equivalence relation between frames, called static equivalence, which we define first. Intuitively, two frames are statically equivalent if no ‘test’ $M = N$ can tell them apart, where M and N have variables that are substituted from the frame. Formally:

Definition 9 (Static equivalence) *Two closed frames φ and ψ are statically equivalent, denoted $\varphi \approx_s \psi$, if $\text{dom}(\varphi) = \text{dom}(\psi)$ and there exists a set of names \tilde{n} and substitutions σ, τ such that $\varphi \equiv \nu \tilde{n}.\sigma$ and $\psi \equiv \nu \tilde{n}.\tau$ and for all terms M, N such that $\tilde{n} \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$, we have $M\sigma =_E N\sigma$ holds if and only if $M\tau =_E N\tau$ holds. We say two closed extended processes A, B are statically equivalent, and write $A \approx_s B$, when their frames are statically equivalent; that is, $\varphi(A) \approx_s \varphi(B)$.*

The relation is called *static* equivalence because it only examines the current state of the processes (as represented by their frames), and not the processes’ dynamic behaviour (that is, the ways in which they may execute in the future). Thus, two processes are statically equivalent if they cannot be distinguished on the basis of their output so far. Static equivalence captures the static part of observational equivalence. More precisely, observational equivalence and static equivalence coincide on frames. Static equivalence can straightforwardly be shown to be closed under structural equivalence, internal reduction and application of closing evaluation contexts.

Example 10 (Static equivalence) *In the following examples we assume the signature $\Sigma_S = \{\text{hash}, \text{fst}, \text{snd}, \text{pair}\}$ and the smallest equational theory E_S satisfying the equations $\text{fst}(\text{pair}(x, y)) = x$, $\text{snd}(\text{pair}(x, y)) = y$ over all variables x, y , where as expected hash , fst , snd are unary functions and pair is a binary function.*

- $\nu m.\{m/x\} \approx_s \nu n.\{n/x\}$; trivial, since they are structurally equivalent.
- $\nu m.\{m/x\} \approx_s \nu n.\{\text{hash}(n)/x\}$.
- $\{m/x\} \not\approx_s \{\text{hash}(m)/x\}$. The first one satisfies $x = m$ but the second one does not.
- $\nu k.\{k/x\} \mid \nu s.\{s/y\} \approx_s \nu k.\{\text{hash}(\text{pair}(a, k))/x\} \mid \{\text{hash}(\text{pair}(b, k))/y\}$.
- $\nu k.\{k/x\} \mid \nu s.\{s/y\} \not\approx_s \nu k.\{\{k/x\} \mid \{\text{hash}(k)/y\}\}$, since the second one satisfies $\text{hash}(x) = y$ and the first one does not.
- $\nu s.\{\text{pair}(s, s)/x\} \not\approx_s \nu s.\{s/x\}$, since the first one satisfies $\text{pair}(\text{fst}(x), \text{snd}(x)) = x$ but the second one does not.

As mentioned, static equivalence captures the static part of observational equivalence. The following definition of labelled bisimilarity captures the dynamic part.

Definition 10 (Labelled bisimilarity) *Labelled bisimilarity (\approx_l) is the largest symmetric relation \mathcal{R} on closed extended processes such that $A \mathcal{R} B$ implies:*

1. $A \approx_s B$;
2. if $A \rightarrow A'$, then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' ;

3. if $A \xrightarrow{\alpha} A'$ and $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$; then $B \rightarrow^* \xrightarrow{\alpha} \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' .

Clauses 2 and 3 of this definition correspond to classical notions of bisimilarity [Mil99]. Notice the use of the “largest relation” construction, to allow us to insist that the processes A' and B' are again within the relation. Clause 1 asserts static equivalence at each step of the bisimulation.

Let us now consider the side condition of Clause 3. Recall that there are three possibilities for α , namely $c(M)$, $\bar{c}\langle u \rangle$, and $\nu u.\bar{c}\langle u \rangle$, where u is either a channel name or a variable of base type. In Clause 3, α can have free variables (in the cases $c(M)$ and $\bar{c}\langle u \rangle$), but any such variables must be defined by the frame of A (and since $A \approx_s B$, also by the frame of B). Hence the condition $\text{fv}(\alpha) \subseteq \text{dom}(A)$. The label α can also have bound names (in the case $\nu d.\bar{c}\langle d \rangle$, where d is a channel name), in which case the transition $A \xrightarrow{\alpha} A'$ has the effect of removing the restriction on a bound channel name. That channel must not occur free in B , to avoid confusing the restricted name with the global name; hence the condition $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$. To see this in more detail, consider A, B such that $A \approx_l B$, and suppose $a \in \text{fn}(B) \setminus \text{fn}(A)$. (For example, A is 0 and B is if $a = b$ then $\bar{c}\langle a \rangle$.) Intuitively, we would like $A \mid (\nu a.\bar{c}\langle a \rangle) \approx_l B \mid (\nu a.\bar{c}\langle a \rangle)$. But to achieve that, we need the side condition, for otherwise the transition $A \mid (\nu a.\bar{c}\langle a \rangle) \xrightarrow{\nu a.\bar{c}\langle a \rangle} A$ would have to be matched on the right hand side by $B \mid (\nu a.\bar{c}\langle a \rangle) \xrightarrow{\nu a.\bar{c}\langle a \rangle} B$, which is false.

Working with labelled bisimilarity. To show labelled bisimilarity of A and B , it is necessary to find a symmetric relation \mathcal{R} that satisfies the conditions contained in Definition 10, such that $A \mathcal{R} B$. Note that the \mathcal{R} we find is not required to be the largest one. That is because the set of relations that satisfies the conditions contained in Definition 10 is closed under union. Therefore, the largest one is the union of them all. *Any* relation satisfying the conditions of Definition 10 can serve as a witness to show that two processes are labelled bisimilar. Note also that, although labelled bisimilarity is easily seen to be an equivalence relation, the relation \mathcal{R} is merely required to be symmetric.

Example 11 (Labelled bisimilarity) *We prove that for all closed processes P_1, P_2, C and names \tilde{n} such that $c \in \tilde{n}$, we have $\nu \tilde{n}.\langle \bar{c}\langle a \rangle.P_1 \mid c(x).P_2 \rangle \mid C \approx_l \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle \mid C$. Intuitively, the equivalence holds because the only choice that the left side has, which the right side does not, is an internal reduction (private channel communication). Let $A \mathcal{R} B$ hold if there exists a closed extended process C' such that one of the following holds:*

- $A \equiv \nu \tilde{n}.\langle \bar{c}\langle a \rangle.P_1 \mid c(x).P_2 \rangle \mid C'$ and $B \equiv \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle \mid C'$; or
- $B \equiv \nu \tilde{n}.\langle \bar{c}\langle a \rangle.P_1 \mid c(x).P_2 \rangle \mid C'$ and $A \equiv \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle \mid C'$; or
- $A \equiv B$.

One may show that \mathcal{R} satisfies the conditions contained in Definition 10. First, \mathcal{R} is easily seen to be symmetric. For the next part, suppose $A \mathcal{R} B$.

1. *We show $A \approx_s B$. There are three cases.*

- (a) $A \equiv \nu \tilde{n}.\langle \bar{c}\langle a \rangle.P_1 \mid c(x).P_2 \rangle \mid C'$ and $B \equiv \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle \mid C'$. Since $\nu \tilde{n}.\langle \bar{c}\langle a \rangle.P_1 \mid c(x).P_2 \rangle \approx_s \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle$, we have $A \approx_s B$.
- (b) $B \equiv \nu \tilde{n}.\langle \bar{c}\langle a \rangle.P_1 \mid c(x).P_2 \rangle \mid C'$ and $A \equiv \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle \mid C'$. In this case, a similar argument applies.
- (c) $A \equiv B$ and hence we trivially have $A \approx_s B$.

2. *Now suppose $A \rightarrow A'$. There are again three cases to consider.*

- (a) $A \equiv \nu \tilde{n}.\langle \bar{c}\langle a \rangle.P_1 \mid c(x).P_2 \rangle \mid C'$ and $B \equiv \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle \mid C'$. If $A' \equiv \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle \mid C'$, then let $B' \triangleq B$. Otherwise, $A' \equiv \nu \tilde{n}.\langle \bar{c}\langle a \rangle.P_1 \mid c(x).P_2 \rangle \mid C''$ for some C'' , in which case let $B' \triangleq \nu \tilde{n}.\langle P_1 \mid P_2\{a/x\} \rangle \mid C''$.

- (b) $A \equiv \nu \tilde{n}.(P_1 \mid P_2\{a/x\}) \mid C'$ and $B \equiv \nu \tilde{n}.(\bar{c}(a).P_1 \mid c(x).P_2) \mid C'$. Since $B \rightarrow A$, it follows that $B \rightarrow^* A'$, so let $B' \triangleq A'$.
- (c) $A \equiv B$ and hence the result follows trivially by $B' \triangleq A'$.

In all cases we have $A' \mathcal{R} B'$, as required.

3. Now suppose $A \xrightarrow{\alpha} A'$. Again we have three cases:

- (a) $A \equiv \nu \tilde{n}.(\bar{c}(a).P_1 \mid c(x).P_2) \mid C'$ and $B \equiv \nu \tilde{n}.(P_1 \mid P_2\{a/x\}) \mid C'$. Since only the C' part is able to make a labelled transition, $A' \equiv \nu \tilde{n}.(\bar{c}(a).P_1 \mid c(x).P_2) \mid C''$ for some C'' , in which case let $B' \triangleq \nu \tilde{n}.(P_1 \mid P_2\{a/x\}) \mid C''$, and we have $B \xrightarrow{\alpha} B'$.
- (b) $A \equiv \nu \tilde{n}.(P_1 \mid P_2\{a/x\}) \mid C'$ and $B \equiv \nu \tilde{n}.(\bar{c}(a).P_1 \mid c(x).P_2) \mid C'$. Since $B \rightarrow A$, it follows that $B \rightarrow^* \xrightarrow{\alpha} A'$, so let $B' \triangleq A'$.
- (c) $A \equiv B$ and hence the result is trivial by $B' \triangleq A'$.

In all cases we have $A' \mathcal{R} B'$, as required.

Abadi & Fournet [AF01] state the following useful results, although as far as we are aware fully detailed proofs have not yet been published.

Lemma 1 *Given closed extended processes A, B and a closing evaluation context $C[-]$, we have $A \approx_l B$ implies $C[A] \approx_l C[B]$.*

Theorem 1 *Observational equivalence and labelled bisimilarity coincide: $\approx = \approx_l$.*

The condition stating that active substitutions are of base type (Section 2.1) is crucial to Theorem 1. Without this constraint, one would have the following counterexample: $\nu a.(\{a/x\} \mid x(y).\bar{c}(n)) \not\approx \nu a.\{a/x\}$ (this can be seen using the closing evaluation context $C[-] = \bar{x}(b) \mid -$); whereas $\nu a.(\{a/x\} \mid x(y).\bar{c}(n)) \approx_l \nu a.\{a/x\}$.

4.3 Strong secrecy

Our earlier definition of secrecy required that the secret was not obtained by the adversary. We will now use notions of equivalence to demonstrate a stronger notion of secrecy which states an adversary is unable to distinguish when the secret changes [Aba00, CRZ07]. Intuitively, this means the value of the secret should not effect the observable behaviour of the protocol. Strong secrecy is useful, for example, to capture the adversary's inability to learn any partial information about the secret.

Strong secrecy is also useful to formalise ‘dictionary attacks’, also known as ‘guessing attacks’, in which the attacker tries all the possible values of the secret until he finds the right one. If a secret is chosen from a relatively small set of possible values, then such an attack might be possible; we call such a secret “weak”. Passwords are often examples of weak secrets. For example, a server which requires the user to send her password s , deterministically encrypted with server's public key pk , that is, to send $\{s\}_{pk}^a$, is vulnerable to guessing attacks on s . That is because an attacker in possession of $\{s\}_{pk}^a$ can make a guess s' of the password s , and check the validity of his guess by constructing $\{s'\}_{pk}^a$ and comparing it with $\{s\}_{pk}^a$. Suppose instead the server requires the user to include a random nonce in the encryption, that is, to send $\{\langle s, r \rangle\}_{pk}^a$, where r is a nonce chosen from a large set. In that case the system is not vulnerable to guessing attacks.

This may be formalised by the following definitions [Bau05, DKR08].

Definition 11 *Let $\varphi \equiv \nu n.\varphi'$ be a frame. We say that φ is resistant to guessing attacks against n if, and only if, $\nu n.(\varphi' \mid \{n/x\}) \approx_s \nu n'.\nu n.(\varphi' \mid \{n'/x\})$ where n' is a fresh name and x is a variable such that $x \notin \text{dom}(\varphi)$.*

The definition says that the environment cannot distinguish between a situation in which he possesses the correct value of the secret, from another one in which he has some random value instead. This definition can be extended to the more general case.

Definition 12 (Guessing attacks) *Let A be a process and $n \in \text{bn}(A)$. We say that A is resistant to guessing attacks against n if, for every process B such that $A(\rightarrow^* \xrightarrow{\alpha} \rightarrow^*)^* B$, then we have that $\varphi(B)$ is resistant to guessing attacks against n .*

Example 12 (TPM authentication [CR08]) *We consider a protocol where a client \mathcal{A} shares a secret s with a server \mathcal{B} (the protocol is based on the one used in the Trusted Platform Module [Tru07]). In this example, the terms M and N are commonly used for tuples, so it is convenient to introduce a little syntactic sugar to express them. We write (M_1, \dots, M_n) to mean $\text{pair}(M_1, \text{pair}(M_2, \text{pair}(\dots, \text{pair}(M_n, *) \dots)))$, where $*$ is any constant. We also write $\text{1st}(M)$, $\text{2nd}(M)$, and $\text{3rd}(M)$ to mean $\text{fst}(M)$, $\text{fst}(\text{snd}(M))$, and $\text{fst}(\text{snd}(\text{snd}(M)))$ respectively. Note that (M_1, \dots, M_k) , $\text{1st}(N)$, etc., are not new functions, but merely syntactic sugar for combinations of functions that we already have. The protocol proceeds as follows. The client sends commands to the server, and authenticates them by supplying a MAC of the command values, keyed on s . An additional nonce n is supplied and used in the MAC, to ensure it has high entropy. More precisely, to run the command ‘comm’, \mathcal{A} sends the message $(\text{comm}, n, \text{mac}(s, (\text{comm}, n)))$. \mathcal{B} checks the MAC by reconstructing it, and if it is correct, sends the response ‘resp’. This protocol may be modelled by the following process P , where we assume the signature $\Sigma_T = \Sigma_S \cup \{\text{mac}\}$ and equational theory E_S (recall that Σ_S and E_S were defined in Example 10, and note that binary function mac has no equations).*

$$\begin{aligned} P &\triangleq \nu s.(!P_A \mid !P_B) \\ P_A &\triangleq \nu n.\bar{c}\langle(\text{comm}, n, \text{mac}(s, (\text{comm}, n)))\rangle \\ P_B &\triangleq c(x).\text{if } \text{3rd}(x) = \text{mac}(s, (\text{1st}(x), \text{2nd}(x))) \text{ then } \bar{c}\langle\text{resp}\rangle \end{aligned}$$

P is vulnerable to guessing attacks on s . To see this, we consider the transition

$$P \xrightarrow{\nu x.\bar{c}\langle x \rangle} \nu s.(\nu n.\{(\text{comm}, n, \text{mac}(s, (\text{comm}, n)))/x\} \mid !P_A \mid !P_B)$$

The frame of this latter process is vulnerable to guessing attacks on s , since we have

$$\begin{aligned} &\nu s.\nu n.(\{(\text{comm}, n, \text{mac}(s, (\text{comm}, n)))/x\} \mid \{s/z\}) \\ &\not\approx_s \nu s'.\nu s.\nu n.(\{(\text{comm}, n, \text{mac}(s, (\text{comm}, n)))/x\} \mid \{s'/z\}) \end{aligned}$$

as witnessed by the test

$$\text{3rd}(x) = \text{mac}(z, (\text{1st}(x), \text{2nd}(x))).$$

A comparison of two views of secrecy. The reachability-based secrecy introduced in Section 3.1 requires that the adversary should never learn secret values; the equivalence-based (strong) secrecy introduced in this section states that two executions of a protocol are indistinguishable regardless of the chosen secret values. The formulation of strong secrecy offers a higher degree of security, but reachability-based secrecy has been more widely studied resulting in more automated support. For further discussion on the relationship between these two styles of secrecy the reader is referred to [CRZ07, Aba00], see also Blanchet for a more generic formalisation of strong secrecy [Bla04].

4.4 Vote privacy for electronic voting

Electronic voting systems are being introduced, or trialed, in several countries to provide more efficient voting procedures with an increased level of security. The applied pi calculus can be used


to formalise the properties which an electronic voting scheme should satisfy [KRS10, SRKK10, DKR09b, BHM08]. For example, the property of vote privacy (or ballot secrecy) means that no-one (not even the election officers) can link a particular vote with the voter.

Consider the following protocol, which is a simplification of the one by Fujioka, Okamoto, and Ohta [FOO92]. We assume the signature $\Sigma_F = \{\text{true}, \text{fst}, \text{snd}, \text{pk}, \text{getmsg}, \text{pair}, \text{sign}, \text{checksign}, \text{blind}, \text{unblind}\}$, the arities of which are defined in the usual way; that is, true is a constant; $\text{fst}, \text{snd}, \text{pk}, \text{getmsg}$ are unary functions; and $\text{pair}, \text{sign}, \text{checksign}, \text{blind}, \text{unblind}$ are binary functions. The protocol [FOO92] relies on *blind signatures*; with this cryptographic primitive, an election officer can sign a text without having seen it. The voter first blinds the text and the officer signs it. The voter can then unblind the message to recover the officer's signature on the text. We do not need to consider how this cryptography actually works; we can encode the effect using the equation

$$\text{unblind}(x, \text{sign}(y, \text{blind}(x, z))) = \text{sign}(y, z)$$

which says that if one blinds a text z using a blinding factor x , and then signs it with the private key y , and then unblinds it (again using x), then the result is the text z signed by y in the usual way. Blind signatures are useful in electronic voting because they allow the officer to sign a vote (thus establishing its eligibility for counting) without knowing what the vote is. The equations associated with the remaining functions in Σ_F are defined in the usual way (see Section 2.1.1).

The protocol proceeds as follows. The voter wishing to vote for candidate v creates a nonce n , and blinds the tuple $\langle v, n \rangle$ using a random blinding factor r . She then signs this value and sends it to the election officer. The officer checks the voter's eligibility and also checks she has not already voted. If these checks succeed, then the officer signs the blinded vote-nonce pair and sends it back to the voter. The voter can now unblind the signature, recovering the officer's signature on the vote-nonce pair. Once all voters have obtained the officer's signature in this way, they can anonymously submit the signed vote-nonce pair for counting. The protocol thus proceeds as follows:



The purpose of the value n is to allow the officer to count the signed vote only once. The synchronisation point, denoted *synch*, ensures that every voter receives the officer's signature before any of them submit the unblinded value for counting. It is necessary to ensure vote privacy; without it, traffic analysis might link the voter's signature in the first message with the anonymous submission in the third one. This protocol is inadequate for a real election, but it does satisfy the property of vote privacy.

We formalise this illustrative voting protocol in the applied pi calculus as follows. We assume a public channel c for communication between the voter and the officer, and a channel $vote$ for

the officer to declare the list of counted votes.

$$\begin{aligned}
P &\triangleq \nu sk_1 \dots \nu sk_n. \nu sk_O. \\
&\text{let } pk_1 = \text{pk}(sk_1) \text{ in } \dots \text{let } pk_n = \text{pk}(sk_n) \text{ in} \\
&\text{let } pk_O = \text{pk}(sk_O) \text{ in} \\
&(\bar{c}\langle pk_1 \rangle \mid \dots \mid \bar{c}\langle pk_n \rangle \mid \bar{c}\langle pk_O \rangle \mid P_V\{sk_1/sk_V, v_1/v\} \mid \dots \mid P_V\{sk_n/sk_V, v_n/v\} \mid !P_O \mid S) \\
P_V &\triangleq \nu n. \nu r. \text{let } bvn = \text{blind}(r, \text{pair}(v, n)) \text{ in} \\
&\bar{c}\langle \text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, bvn)) \rangle. \\
&c(x). \text{if } \text{checksign}(pk_O, x) = \text{true} \text{ then} \\
&\text{if } \text{getmsg}(x) = bvn \text{ then} \\
&\textit{synch}. \\
&\bar{c}\langle \text{unblind}(r, x) \rangle \\
P_O &\triangleq c(y). \text{if } \text{checksign}(\text{fst}(y), \text{snd}(y)) = \text{true} \text{ then} \\
&\text{if } \text{Eligible}(\text{fst}(y)) = \text{true} \text{ then} \\
&\bar{c}\langle \text{sign}(sk_O, \text{getmsg}(\text{snd}(y))) \rangle. \\
&c(w). \\
&\text{if } \text{checksign}(pk_O, w) = \text{true} \text{ then} \\
&\text{if } \text{NotSeen}(w) = \text{true} \text{ then} \\
&\overline{\textit{vote}}\langle \text{fst}(\text{getmsg}(w)) \rangle
\end{aligned}$$

In the above process, we assume functions $\text{Eligible}(x)$ and $\text{NotSeen}(x)$ are available to the officer. $\text{Eligible}(x)$ checks whether the voter with public key x is eligible and hasn't already voted; $\text{NotSeen}(x)$ stores x and returns true if it has not seen x before; otherwise it returns false. The purpose of $\text{NotSeen}(x)$ is to ensure that the voter can't use the signed token twice. We also assume a subprocess \textit{synch} which a voter uses to synchronise with other voters. This can be modelled by defining a process S to coordinate the synchronisation. In the case of n voters, the coordination process receives precisely n tokens, and then sends them back. To synchronise with the other voters, a voter sends a token to this synchronisation coordinator and awaits its response. Thus,

$$\begin{aligned}
\textit{synch} &\triangleq \overline{\textit{syn}}\langle * \rangle. \textit{syn}'(o) \\
S &\triangleq \textit{syn}(x_1) \dots \textit{syn}(x_n). \overline{\textit{syn}}'\langle * \rangle \dots \overline{\textit{syn}}'\langle * \rangle
\end{aligned}$$

where $\textit{syn}, \textit{syn}'$ are private channels shared between the voters and S , $*$ is any name and o is a variable.

Now we define the property of vote privacy, and prove it. Since it should hold even if the officer is corrupt, we can suppose that the officer is part of the attacker, and we need not consider its formalisation as P_O above. This means the secrecy of the officer's key is not required for the vote privacy property, so we can treat it as a free name. Moreover, the voters' secret keys are also not required to be kept secret in order to ensure the vote privacy property. They too can be treated as free names. (Of course, these keys *are* required to be secret for other properties.)

Consider two voters \mathcal{A}, \mathcal{B} and two candidates v_a, v_b . Based upon [DKR09b], we formalise vote privacy for two voters with the assertion that the attacker (including the election officers) cannot distinguish between a situation in which \mathcal{A} votes v_a and \mathcal{B} votes v_b , from another one in which \mathcal{A} votes v_b and \mathcal{B} votes v_a . We will write this as the equivalence:

$$\begin{aligned}
\nu \textit{syn}. \nu \textit{syn}' . (P_V\{sk_A/sk_V, v_a/v\} \mid P_V\{sk_B/sk_V, v_b/v\} \mid S) \\
\approx_l \nu \textit{syn}. \nu \textit{syn}' . (P_V\{sk_A/sk_V, v_b/v\} \mid P_V\{sk_B/sk_V, v_a/v\} \mid S)
\end{aligned}$$

To prove this equivalence, let us call the left hand side LHS, and the right hand side RHS. We need to show a relation \mathcal{R} satisfying the requirements of Definition 10 such that $\text{LHS} \mathcal{R} \text{RHS}$. As we have seen, the idea of \mathcal{R} is to relate successors of LHS and RHS that are expected to correspond to each other in the bisimulation. As LHS evolves, the corresponding evolution of RHS is the one that mimics \mathcal{A} moves in LHS with \mathcal{A} moves in RHS, and \mathcal{B} moves in LHS with \mathcal{B} moves in

RHS, up to the synchronisation. After the synchronisation, \mathcal{A} moves in LHS are mimicked by \mathcal{B} moves in RHS, and \mathcal{B} moves in LHS are mimicked by \mathcal{A} moves in RHS. The reason for the ‘swap’ in mimicking is to ensure that the static equivalences will hold; before the synchronisation, the output data produced by \mathcal{A} on both the LHS and RHS are indistinguishable (and similarly for \mathcal{B}). Observe that the output data reveals the keys $\text{pk}(sk_A)$, $\text{pk}(sk_B)$ (matched by $\text{pk}(sk_A)$, $\text{pk}(sk_B)$) and the signed blind votes $\text{sign}(sk_A, \text{blind}(r, \text{pair}(v_a, n)))$, $\text{sign}(sk_B, \text{blind}(r', \text{pair}(v_b, n')))$ (matched by $\text{sign}(sk_A, \text{blind}(r, \text{pair}(v_b, n)))$, $\text{sign}(sk_B, \text{blind}(r', \text{pair}(v_a, n')))$) where names n, n', r, r' are under restriction in both LHS and RHS. Indistinguishability between the LHS and RHS with respect to the signed blind votes is due to the properties of blinding. After the synchronisation \mathcal{A} will reveal v_a on the LHS and v_b on the RHS (similarly \mathcal{B} will reveal v_b on the LHS and v_a on the RHS). Hence after synchronisation the actions mimicked are swapped.

To define this formally, we first define some partial evolutions of P_V , as follows.

$$\begin{aligned}
P_1 &\triangleq \nu n. \nu r. \\
&\quad \bar{c}(\text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, \text{blind}(r, \text{pair}(v, n))))) \\
&\quad c(x). \text{if } \text{checksign}(pk_O, x) = \text{true} \text{ then} \\
&\quad \text{if } \text{getmsg}(x) = \text{blind}(r, \text{pair}(v, n)) \text{ then} \\
&\quad \overline{\text{syn}}\langle * \rangle. \text{syn}'(o). \\
&\quad \bar{c}(\text{unblind}(r, x)) \\
P_2 &\triangleq \nu n. \nu r. (\{\text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, \text{blind}(r, \text{pair}(v, n))))/y\} \mid \\
&\quad c(x). \text{if } \text{checksign}(pk_O, x) = \text{true} \text{ then} \\
&\quad \text{if } \text{getmsg}(x) = \text{blind}(r, \text{pair}(v, n)) \text{ then} \\
&\quad \overline{\text{syn}}\langle * \rangle. \text{syn}'(o). \\
&\quad \bar{c}(\text{unblind}(r, x))) \\
P_3 &\triangleq \nu n. \nu r. (\{\text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, \text{blind}(r, \text{pair}(v, n))))/y\} \mid \\
&\quad \text{if } \text{checksign}(pk_O, m) = \text{true} \text{ then} \\
&\quad \text{if } \text{getmsg}(m) = \text{blind}(r, \text{pair}(v, n)) \text{ then} \\
&\quad \overline{\text{syn}}\langle * \rangle. \text{syn}'(o). \\
&\quad \bar{c}(\text{unblind}(r, m))) \\
P_4 &\triangleq \nu n. \nu r. (\{\text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, \text{blind}(r, \text{pair}(v, n))))/y\} \mid \\
&\quad \text{if } \text{getmsg}(m) = \text{blind}(r, \text{pair}(v, n)) \text{ then} \\
&\quad \overline{\text{syn}}\langle * \rangle. \text{syn}'(o). \\
&\quad \bar{c}(\text{unblind}(r, m))) \\
P_5 &\triangleq \nu n. \nu r. (\{\text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, \text{blind}(r, \text{pair}(v, n))))/y\} \mid \\
&\quad \overline{\text{syn}}\langle * \rangle. \text{syn}'(o). \\
&\quad \bar{c}(\text{unblind}(r, \text{sign}(sk_O, \text{blind}(r, \text{pair}(v, n))))) \\
P_6 &\triangleq \nu n. \nu r. (\{\text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, \text{blind}(r, \text{pair}(v, n))))/y\} \mid \\
&\quad \text{syn}'(o). \\
&\quad \bar{c}(\text{unblind}(r, \text{sign}(sk_O, \text{blind}(r, \text{pair}(v, n))))) \\
P_7 &\triangleq \nu n. \nu r. (\{\text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, \text{blind}(r, \text{pair}(v, n))))/y\} \mid \\
&\quad \bar{c}(\text{unblind}(r, \text{sign}(sk_O, \text{blind}(r, \text{pair}(v, n))))) \\
P_8 &\triangleq \nu n. \nu r. (\{\text{pair}(\text{pk}(sk_V), \text{sign}(sk_V, \text{blind}(r, \text{pair}(v, n))))/y\} \mid \\
&\quad \{\text{unblind}(r, \text{sign}(sk_O, \text{blind}(r, \text{pair}(v, n)))/u\})
\end{aligned}$$

These are abbreviations that will be useful in our definition of the simulation relation. We define the relation \mathcal{R} as follows. Given closed extended processes X and Y , $X \mathcal{R} Y$ and $Y \mathcal{R} X$ both hold if

- there exist integers i, j , variables w, z and terms M, N with $1 \leq i, j \leq 6$ and

$$\begin{aligned} X &\equiv P_i\{sk_A/sk_V, v_a/v, w/y, M/m\} \mid P_j\{sk_B/sk_V, v_b/v, z/y, N/m\} \mid S_{i,j}, \\ Y &\equiv P_i\{sk_A/sk_V, v_b/v, w/y, M/m\} \mid P_j\{sk_B/sk_V, v_a/v, z/y, N/m\} \mid S_{i,j}; \end{aligned}$$

and if $i = 4$ then $\text{checksign}(pk_O, M) = \text{true}$, and if $j = 4$ then $\text{checksign}(pk_O, N) = \text{true}$; or

- there exist integers i, j , and variables s, t, w, z with $6 \leq i, j \leq 8$ and

$$\begin{aligned} X &\equiv P_i\{sk_A/sk_V, v_a/v, w/y, s/u\} \mid P_j\{sk_B/sk_V, v_b/v, z/y, t/u\} \mid S_{i,j}, \\ Y &\equiv P_j\{sk_A/sk_V, v_b/v, w/y, t/u\} \mid P_i\{sk_B/sk_V, v_a/v, z/y, s/u\} \mid S_{i,j}. \end{aligned}$$

Notice that the substitutions above are syntactic, and recall that in an active substitution, the variable in the domain is considered to occur free. Thus, $(A \mid \{M/y\})\{x/y\} = A\{x/y\} \mid \{M\{x/y\}/x\}$; that is, the domain variable changed from y to x .

In the definition of \mathcal{R} , $S_{i,j}$ is whatever part of S is left to execute. More precisely,

$$S_{i,j} \triangleq \begin{cases} \text{syn}(x)\text{syn}(x).\overline{\text{syn}}\langle * \rangle \overline{\text{syn}}\langle * \rangle & \text{if } i, j < 6 \\ \text{syn}(x).\overline{\text{syn}}\langle * \rangle \overline{\text{syn}}\langle * \rangle & \text{if } i = 6, j < 6 \text{ or } i < 6, j = 6 \\ \overline{\text{syn}}\langle * \rangle \overline{\text{syn}}\langle * \rangle & \text{if } i, j = 6 \\ \overline{\text{syn}}\langle * \rangle & \text{if } i = 6, j = 7 \text{ or } i = 7, j = 6 \\ 0 & \text{if } i, j \geq 7 \end{cases}$$

We present a proof sketch that \mathcal{R} satisfies the requirements of Definition 10.

Static equivalence. Consider, for example:

$$\begin{aligned} P_1\{sk_A/sk_V, v_a/v\} \mid P_2\{sk_B/sk_V, v_b/v, z/y\} \mid S_{1,2} \\ \mathcal{R} \quad P_1\{sk_A/sk_V, v_b/v\} \mid P_2\{sk_B/sk_V, v_a/v, z/y\} \mid S_{1,2} \end{aligned}$$

The frame of the process on the left is $\nu n_1.\nu r_1.\{\text{pair}(\text{pk}(sk_B), \text{sign}(sk_B, \text{blind}(r_1, \text{pair}(v_b, n_1))))/z\}$ while the one on the right is $\nu n_1.\nu r_1.\{\text{pair}(\text{pk}(sk_B), \text{sign}(sk_B, \text{blind}(r_1, \text{pair}(v_a, n_1))))/z\}$. These differ only by the subterm v_b on the left (v_a on the right). Static equivalence holds because of the blinding. Consider

$$\begin{aligned} P_7\{sk_A/sk_V, v_a/v, w/y\} \mid P_8\{sk_B/sk_V, v_b/v, z/y, t/u\} \mid S_{7,8} \\ \mathcal{R} \quad P_7\{sk_B/sk_V, v_a/v, z/y\} \mid P_8\{sk_A/sk_V, v_b/v, w/y, t/u\} \mid S_{7,8} \end{aligned}$$

The frame of the process on the left is structurally equivalent to

$$\begin{aligned} \nu n_1.\nu r_1.\{\text{pair}(\text{pk}(sk_A), \text{sign}(sk_A, \text{blind}(r_1, \text{pair}(v_a, n_1))))/w\} \mid \\ \nu n_2.\nu r_2.\{\{\text{pair}(\text{pk}(sk_B), \text{sign}(sk_B, \text{blind}(r_2, \text{pair}(v_b, n_2))))/z\} \mid \\ \{\text{unblind}(r_2, \text{sign}(sk_O, \text{blind}(r_2, \text{pair}(v_b, n_2))))/t\}\} \end{aligned}$$

while the frame on the right is structurally equivalent to

$$\begin{aligned} \nu n_1.\nu r_1.\{\text{pair}(\text{pk}(sk_B), \text{sign}(sk_B, \text{blind}(r_1, \text{pair}(v_a, n_1))))/z\} \mid \\ \nu n_2.\nu r_2.\{\{\text{pair}(\text{pk}(sk_A), \text{sign}(sk_A, \text{blind}(r_2, \text{pair}(v_b, n_2))))/w\} \mid \\ \{\text{unblind}(r_2, \text{sign}(sk_O, \text{blind}(r_2, \text{pair}(v_b, n_2))))/t\}\} \end{aligned}$$

These frames are statically equivalent.

Unlabelled transitions. Suppose, for example, $A \rightarrow A'$ and

$$\begin{aligned} A &\equiv P_3\{sk_A/sk_V, v_a/v, w/y, M/m\} \mid P_4\{sk_B/sk_V, v_b/v, z/y, N/m\} \mid S_{3,4} \\ A' &\equiv P_3\{sk_A/sk_V, v_a/v, w/y, M/m\} \mid P_5\{sk_B/sk_V, v_b/v, z/y\} \mid S_{3,5} \end{aligned}$$

Then, by \mathcal{R} ,

$$B \equiv P_3\{sk_A/sk_V, v_b/v, w/y, M/m\} \mid P_4\{sk_B/sk_V, v_a/v, z/y, N/m\} \mid S_{3,4}$$

and $\text{checksign}(pk_O, N) = \text{true}$. Since the transition $A \rightarrow A'$ was possible, we have $\text{getmsg}(N) = \text{blind}(r, \text{pair}(v_b, n))$, where r, n are the values used inside the occurrence of P_4 . It follows that we have $N = \text{sign}(sk_O, \text{getmsg}(\text{snd}(z)))$. Therefore, the transition $B \rightarrow B'$ is possible, where

$$B' \triangleq P_3\{sk_A/sk_V, v_b/v, w/y, M/m\} \mid P_5\{sk_B/sk_V, v_a/v, z/y\} \mid S_{3,5},$$

and we have $A' \mathcal{R} B'$.

Labelled transitions. Suppose, for example, $A \xrightarrow{\nu w.\bar{c}\langle w \rangle} A'$ and

$$\begin{aligned} A &\equiv P_1\{sk_A/sk_V, v_a/v\} \mid P_2\{sk_B/sk_V, v_b/v, z/y\} \mid S_{1,2} \\ A' &\equiv P_2\{sk_A/sk_V, v_a/v, w/y\} \mid P_2\{sk_B/sk_V, v_b/v, z/y\} \mid S_{2,2} \end{aligned}$$

Then, by \mathcal{R} ,

$$B \equiv P_1\{sk_A/sk_V, v_b/v\} \mid P_2\{sk_B/sk_V, v_a/v, z/y\} \mid S_{1,2}$$

and so we set

$$B' \triangleq P_2\{sk_A/sk_V, v_b/v, w/y\} \mid P_2\{sk_B/sk_V, v_a/v, z/y\} \mid S_{2,2}.$$

We have $A' \mathcal{R} B'$.

Suppose, for example, $A \xrightarrow{c(T)} A'$ and

$$\begin{aligned} A &\equiv P_1\{sk_A/sk_V, v_a/v\} \mid P_2\{sk_B/sk_V, v_b/v, z/y\} \mid S_{1,2} \\ A' &\equiv P_1\{sk_A/sk_V, v_a/v\} \mid P_3\{sk_B/sk_V, v_b/v, z/y, T/m\} \mid S_{1,3} \end{aligned}$$

Then, by \mathcal{R} ,

$$B \equiv P_1\{sk_A/sk_V, v_b/v\} \mid P_2\{sk_B/sk_V, v_a/v, z/y\} \mid S_{1,2}$$

and so we set

$$B' \triangleq P_1\{sk_A/sk_V, v_b/v\} \mid P_3\{sk_B/sk_V, v_a/v, z/y, T/m\} \mid S_{1,3}$$

We again have $A' \mathcal{R} B'$.

This concludes our proof sketch of the labelled bisimilarity that expresses the privacy property for two voters. To complete this proof formally, one would have to consider all the cases involved in showing that R satisfies the requirements of Definition 10.

5 Outlook

This chapter studies the applied pi calculus and demonstrates its applicability for analysing secrecy, correspondence and observational equivalence properties in the context of security protocols. In this section we briefly refer to some on-going research about the calculus. Our discussion is neither exhaustive nor complete.

The ability to reason with security properties is largely dependent on being able to reason with the equational theory. Abadi & Cortier [AC04] have studied decidability of secrecy and static equivalence for a large class of equational theories in the presence of a passive adversary. Subsequent work by Baudet, Cortier & Delaune [BCD09] and Ciobăcă, Delaune & Kremer [CDK09] introduces automated tools for analysing static equivalence with respect to convergent equational theories.

For an active adversary, reachability properties, in particular secrecy, have been studied in [Bla01, AB02] with updated and extended versions [AB05a, Bla10]. Strong secrecy has been considered in [Bla04]. Correspondence properties have been presented in the context of authenticity [Bla02], and subsequently extended upon and revised in [Bla09]. Proofs of observational equivalence have been studied with respect to processes that differ only by the choice of the terms that they contain, using the notion of uniformity [BAF05, BAF08]. But this notion of uniformity is often too strict; for example, the voting process discussed in this chapter does not satisfy it. A practical approach to overcoming the problem is introduced [DRS08], [Smy11, Chapter 5] (the technique defined by [Smy11, Chapter 5] has been implemented as a tool called ProSwapper [KSR10]) and subsequently used [BHM08, DDS10]. As an alternative to uniformity as a proof technique, a symbolic version of the applied pi calculus is introduced in [DKR09a]; by treating inputs symbolically, it avoids potentially infinite branching of processes due to inputs from the environment. Cortier & Delaune [CD09] have shown that observational equivalence coincides with trace equivalence for determinate processes; and based on the symbolic semantics, this yields a decision procedure for observational equivalence of finite determinate processes without replication.

Acknowledgements

We are very grateful to Myrto Arapinis and the editors of [RS10] for careful reading of a draft of this work. In addition, we are grateful to Jiangshan Yu and Takahiro Kubota for highlighting mistakes in an earlier version of this document. This work has been partly supported by the EPSRC projects: *UbiVal: Fundamental Approaches to Validation of Ubiquitous Computing Applications and Infrastructures* (EP/D076625/2); and *Verifying Interoperability Requirements in Pervasive Systems* (EP/F033540/1). This work was partly done while Ben Smyth was at CNRS, Département d'Informatique, École Normale Supérieure, Paris, France with support from the *Direction Générale pour l'Armement*.

References

- [AB02] Martín Abadi and Bruno Blanchet. Analyzing Security Protocols with Secrecy Types and Logic Programs. In *POPL'02: Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 33–44. ACM Press, 2002.
- [AB05a] Martín Abadi and Bruno Blanchet. Analyzing Security Protocols with Secrecy Types and Logic Programs. *Journal of the ACM*, 52(1):102–146, 2005.
- [AB05b] Martín Abadi and Bruno Blanchet. Computer-Assisted Verification of a Protocol for Certified Email. *Science of Computer Programming*, 58(1–2):3–27, 2005. Special issue SAS'03.
- [Aba00] Martín Abadi. Security Protocols and their Properties. In Friedrich L. Bauer and Ralf Steinbrüggen, editors, *Foundations of Secure Computation*, NATO Science Series, pages 39–60. IOS Press, 2000.
- [ABF07] Martín Abadi, Bruno Blanchet, and Cédric Fournet. Just Fast Keying in the Pi Calculus. *ACM Transactions on Information and System Security*, 10(3):1–59, 2007.
- [AC04] Martín Abadi and Véronique Cortier. Deciding knowledge in security protocols under equational theories. In *ICALP'04: Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, volume 3142 of *LNCS*, pages 46–58. Springer, 2004.
- [AC06] Martín Abadi and Véronique Cortier. Deciding knowledge in security protocols under equational theories. *Theoretical Computer Science*, 367(1–2):2–32, 2006.

- [AF01] Martín Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In *POPL'01: Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 104–115. ACM Press, 2001.
- [AF06] Martín Abadi and Cédric Fournet. Private email communication, 24th May 2006.
- [AG97] Martín Abadi and Andrew D. Gordon. A calculus for cryptographic protocols: the spi calculus. In *CCS'97: Proceedings of the 4th ACM conference on Computer and Communications Security*, pages 36–47. ACM Press, 1997.
- [AN96] Martín Abadi and Roger Needham. Prudent engineering practice for cryptographic protocols. *IEEE Transactions on Software Engineering*, 22(1):6–15, January 1996.
- [BAF05] Bruno Blanchet, Martín Abadi, and Cédric Fournet. Automated verification of selected equivalences for security protocols. In *LICS'05: Proceedings of the 20th IEEE Symposium on Logic in Computer Science*, pages 331–340. IEEE Computer Society, 2005.
- [BAF08] 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, 2008.
- [Bau05] Mathieu Baudet. Deciding security of protocols against off-line guessing attacks. In *CCS'05: Proceedings of the 12th Conference on Computer and Communications Security*, pages 16–25. ACM Press, 2005.
- [BC08] Bruno Blanchet and Avik Chaudhuri. Automated Formal Analysis of a Protocol for Secure File Sharing on Untrusted Storage. In *S&P'08: Proceedings of the 29th IEEE Symposium on Security and Privacy*, pages 417–431. IEEE Computer Society, 2008.
- [BCD09] Mathieu Baudet, Véronique Cortier, and Stéphanie Delaune. YAPA: A generic tool for computing intruder knowledge. In *RTA'09: Proceedings of the 20th International Conference on Rewriting Techniques and Applications*, volume 5595 of *LNCS*, pages 148–163. Springer, 2009.
- [BFGT06] Karthikeyan Bhargavan, Cedric Fournet, Andrew D. Gordon, and Stephen Tse. Verified Interoperable Implementations of Security Protocols. In *CSFW'06: Proceedings of the IEEE Computer Security Foundations Workshop*, pages 139–152. IEEE Computer Society, 2006.
- [BHM08] Michael Backes, Catalin Hritcu, and Matteo Maffei. Automated Verification of Remote Electronic Voting Protocols in the Applied Pi-Calculus. In *CSF'08: Proceedings of the 21st IEEE Computer Security Foundations Symposium*, pages 195–209. IEEE Computer Society, 2008.
- [BJPV09] Jesper Bengtson, Magnus Johansson, Joachim Parrow, and Björn Victor. Psi-calculi: Mobile Processes, Nominal Data, and Logic. In *LICS '09: Proceedings of the 24th Annual IEEE Symposium on Logic In Computer Science*, pages 39–48. IEEE Computer Society, 2009.
- [Bla01] Bruno Blanchet. An efficient cryptographic protocol verifier based on Prolog rules. In *CSFW'01: Proceedings of the 14th IEEE Computer Security Foundations Workshop*, pages 82–96. IEEE Computer Society, 2001.
- [Bla02] Bruno Blanchet. From Secrecy to Authenticity in Security Protocols. In *SAS'02: Proceedings of the 9th International Static Analysis Symposium*, volume 2477 of *LNCS*, pages 342–359. Springer, 2002.

- [Bla04] Bruno Blanchet. Automatic Proof of Strong Secrecy for Security Protocols. In *S&P'04: Proceedings of the 25th IEEE Symposium on Security and Privacy*, pages 86–100. IEEE Computer Society, 2004.
- [Bla09] Bruno Blanchet. Automatic Verification of Correspondences for Security Protocols. *Journal of Computer Security*, 17(4):363–434, 2009.
- [Bla10] Bruno Blanchet. Using Horn clauses for analyzing protocols. In V Cortier and S Kremer, editor, *Formal Models and Techniques for Analyzing Security Protocols*, chapter 5. IOS Press, 2010. To appear.
- [BMU08] Michael Backes, Matteo Maffei, and Dominique Unruh. Zero-Knowledge in the Applied Pi-calculus and Automated Verification of the Direct Anonymous Attestation Protocol. In *S&P'08: Proceedings of the 29th IEEE Symposium on Security and Privacy*, pages 202–215. IEEE Computer Society, 2008.
- [BS10] Bruno Blanchet and Ben Smyth. ProVerif: Automatic Cryptographic Protocol Verifier User Manual & Tutorial. Available from <http://www.proverif.ens.fr/>, 2010.
- [CD09] Véronique Cortier and Stéphanie Delaune. A method for proving observational equivalence. In *CSF'09: Proceedings of the 22nd IEEE Computer Security Foundations Symposium*, pages 266–276. IEEE Computer Society, 2009.
- [CDK09] Ştefan Ciobăcă, Stéphanie Delaune, and Steve Kremer. Computing knowledge in security protocols under convergent equational theories. In *CADE'09: Proceedings of the 22nd International Conference on Automated Deduction*, Lecture Notes in Artificial Intelligence, pages 355–370. Springer, 2009.
- [CR08] Liqun Chen and Mark D. Ryan. Offline dictionary attack on TCG TPM weak authorisation data, and solution. In David Gawrock, Helmut Reimer, Ahmad-Reza Sadeghi, and Claire Vishik, editors, *Future of Trust in Computing*. Vieweg & Teubner, 2008.
- [CR09] Liqun Chen and Mark D. Ryan. Attack, solution and verification for shared authorisation data in TCG TPM. In *FAST'09: Proceedings of the 6th Formal Aspects in Security and Trust*, LNCS. Springer, 2009.
- [CRZ07] Véronique Cortier, Michaël Rusinowitch, and Eugen Zălinescu. Relating two standard notions of secrecy. *Logical Methods in Computer Science*, 3(3), 2007.
- [DDS10] Morten Dahl, Stéphanie Delaune, and Graham Steel. Formal Analysis of Privacy for Vehicular Mix-Zones. In *ESORICS'10: Proceedings of the 15th European Symposium on Research in Computer Security*, volume 6345 of LNCS, pages 55–70. Springer, 2010.
- [DKR08] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. Composition of password-based protocols. In *CSF'08: Proceedings of the 21st IEEE Computer Security Foundations Symposium*, pages 239–251, Pittsburgh, PA, USA, 2008. IEEE Computer Society.
- [DKR09a] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. Symbolic bisimulation for the applied pi calculus. *Journal of Computer Security*, 2009.
- [DKR09b] Stéphanie Delaune, Steve Kremer, and Mark D. Ryan. Verifying Privacy-type Properties of Electronic Voting Protocols. *Journal of Computer Security*, 2009.
- [DRS08] Stéphanie Delaune, Mark D. Ryan, and Ben Smyth. Automatic verification of privacy properties in the applied pi-calculus. In *IFIPTM'08: Proceedings of the 2nd Joint iTrust and PST Conferences on Privacy, Trust Management and Security*, volume 263 of *IFIP Conference Proceedings*, pages 263–278. Springer, 2008.

- [DY83] Danny Dolev and Andrew C. Yao. On the security of public key protocols. *Information Theory*, 29:198–208, 1983.
- [FOO92] Atsushi Fujioka, Tatsuaki Okamoto, and Kazui Ohta. A practical secret voting scheme for large scale elections. In *AUSCRYPT'92: Advances in Cryptology*, volume 718 of *LNCS*, pages 244–251. Springer, 1992.
- [KR05] Steve Kremer and Mark D. Ryan. Analysis of an Electronic Voting Protocol in the Applied Pi-Calculus. In *ESOP'05: Proceedings of the 14th European Symposium On Programming*, volume 3444 of *LNCS*, pages 186–200. Springer, 2005.
- [KRS10] Steve Kremer, Mark D. Ryan, and Ben Smyth. Election verifiability in electronic voting protocols. In *ESORICS'10: Proceedings of the 15th European Symposium on Research in Computer Security*, volume 6345 of *LNCS*, pages 389–404. Springer, 2010. An extended version of this paper appears in [Smy11, Chapter 3].
- [KSR10] Petr Klus, Ben Smyth, and Mark D. Ryan. ProSwapper: Improved equivalence verifier for ProVerif. <http://www.bensmyth.com/proswapper.php>, 2010.
- [KT09] Ralf Küsters and Tomasz Truderung. An Epistemic Approach to Coercion-Resistance for Electronic Voting Protocols. In *S&P'09: Proceedings of the 30th IEEE Symposium on Security and Privacy*, pages 251–266. IEEE Computer Society, 2009.
- [Low97] Gavin Lowe. A hierarchy of authentication specifications. In *CSFW'97: Proceedings of the 10th IEEE workshop on Computer Security Foundations*, pages 31–43. IEEE Computer Society, 1997.
- [MGR09] Aybek Mukhamedov, Andrew D. Gordon, and Mark D. Ryan. Towards a Verified Reference Implementation of the Trusted Platform Module. In *Proceedings of the 17th International Workshop on Security Protocols*, LNCS. Springer, 2009.
- [Mil99] Robin Milner. *Communicating and mobile systems: the π -calculus*. Cambridge University Press, 1999.
- [RS10] Mark D. Ryan and Ben Smyth. Applied pi calculus. In V Cortier and S Kremer, editor, *Formal Models and Techniques for Analyzing Security Protocols*, chapter 6. IOS Press, 2010.
- [SGC07] D. Syme, A. Granicz, and A. Cisternino. *Expert F#*. Apress, 2007.
- [Smy11] Ben Smyth. *Formal verification of cryptographic protocols with automated reasoning*. PhD thesis, School of Computer Science, University of Birmingham, 2011.
- [SRC07] Ben Smyth, Mark D. Ryan, and Liqun Chen. Direct Anonymous Attestation (DAA): Ensuring privacy with corrupt administrators. In *ESAS'07: Proceedings of the 4th European Workshop on Security and Privacy in Ad hoc and Sensor Networks*, volume 4572 of *LNCS*, pages 218–231, 2007. An extended version of this paper appears in [Smy11, Chapter 4].
- [SRKK10] Ben Smyth, Mark D. Ryan, Steve Kremer, and Mounira Kourjeh. Towards automatic analysis of election verifiability properties. In *ARSPA-WITS'10: Proceedings of the Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security*, LNCS. Springer, 2010.
- [Tru07] Trusted Computing Group. TPM Specification version 1.2. Parts 1–3., 2007. Available from <http://www.trustedcomputinggroup.org/specs/TPM/>.