

Causality-based Abstraction of Multiplicity in Security Protocols

Michael Backes Matteo Maffei
Saarland University, Saarbrücken, Germany
{backes,maffei}@cs.uni-sb.de

Agostino Cortesi
Ca' Foscari University, Venice, Italy
cortesi@dsi.unive.it

Abstract

This paper presents a novel technique for analyzing security protocols based on an abstraction of the program semantics. This technique is based on a novel structure called causal graph which captures the causality among program events within a finite graph. A core property of causal graphs is that they abstract away from the multiplicity of protocol sessions, hence constituting a concise tool for reasoning about an even infinite number of concurrent protocol sessions; deciding security only requires a traversal of the causal graph, thus yielding a decidable, and typically very efficient, approach for security protocol analysis. Additionally, causal graphs allow for dealing with different security properties such as secrecy and authenticity in a uniform manner. Both the construction of the causal graph from a given protocol description and the analysis have been fully automated and tested on several example protocols from the literature.

1 Introduction

Proofs of security protocols are known to be error-prone and, owing to the distributed-system aspects of multiple interleaved protocol runs, awkward to make for humans. In fact, vulnerabilities have accompanied the design of such protocols ever since early authentication protocols like Needham-Schroeder [15, 29], over carefully designed de-facto standards like SSL and PKCS [31, 7], up to current widely deployed products like Microsoft Passport [20]. Formal methods, and in particular language-based techniques, have proved to constitute important tools for dealing with such flaws, by helping both to securely design and to analyze security protocols. A central intricacy that these approaches have to tackle is how to concisely treat the potentially unbounded number of concurrent protocol sessions. Techniques based on state-space exploration rely on the inspection of (abstractions of) all the unbounded sequences of messages exchanged by honest principals, and although

such approaches often allow for a careful detection and reproduction of successful attacks, they are inherently constrained by only taking a finite number of sessions into account, or by restricting the analysis to certain classes of protocols, or by giving up guaranteed termination. In contrast to these approaches, static analysis techniques based on type systems rely on the identification of some syntactic patterns that suffice for guaranteeing the property of interest. Furthermore, type systems enjoy guaranteed termination since they work on the syntax of the protocol. The drawback of these techniques is that for analyzing different security properties one has to identify multiple different patterns and, consequently, to develop different type systems. Furthermore, the aforementioned patterns are sometimes restrictive in that, e.g., ensuring freshness of authentication requests presupposes that every authentication session is based on a different nonce (an unguessable random string) that is used only in a single protocol session and then discarded [23, 12]. This rules out some interesting protocols relying on composing multiple authentication sessions based on the same nonce, e.g., see [26].

This paper presents a novel technique for analyzing security protocols based on abstract interpretation of program semantics [13]. In particular, this technique abstracts away from the multiplicity of protocol sessions while still preserving the causality among program events. More precisely, the unbounded number of execution traces of a protocol generated by running an unbounded number of concurrent protocol sessions is concisely abstracted into a novel structure that we call *causal graph*, i.e., a finite graph in which nodes represent process events and edges express the causality among events. Interestingly, causal graphs allow for soundly characterizing which terms abstract messages generated in the same protocol session and which ones may instead abstract messages generated in different protocol sessions; this information turns out to be crucial in order to determine the safety of protocol specifications. In contrast to state-space exploration techniques, our analysis enjoys guaranteed termination. Furthermore, deciding security requires only a traversal of the causal graph, and causal graphs

turn out to be of decent size for commonly analyzed security protocols. Our work improves on existing type-based approaches in that we do not restrict the class of analyzable protocols to the ones adhering to some specific syntactic patterns; this is possible since our analysis is based on an abstraction of the program semantics. Finally, our approach allows for a uniform treatment of different security properties such as secrecy and two different variants of authenticity, namely non-injective agreement and agreement [28].

We can summarize our technique as follows. We first specify the protocol in a dialect of the spi-calculus [4], a process calculus for modelling cryptographic protocols. We then construct the corresponding causal graph, which is proven to be unique for every protocol. The causal graph then allows for analyzing the intended security properties. Since causal graphs are of finite size, the analysis is assured to terminate. Finally, the safety of the causal graph implies the safety of the protocol and, more precisely, the safety of the possibly unbounded number of protocol execution traces. As usual for static analysis and due to the undecidability of the original problem, failures in the verification may be caused by either a flaw in the protocol or a non-sufficient precision of the analysis ruling out safe protocols.

Finally, the analysis is amenable to full automation. We have implemented a tool for automating the analysis, and we have applied the tool to a number of common protocols [18]: the analyses terminated within a few seconds and provided safety proofs for the correct versions of the protocols while failing to validate the flawed versions. Remarkably, attacks are often easily derivable by an inspection of the causal graph. The only human effort required is to capture the protocol in the dialect of the spi-calculus which is often straightforwardly derivable from the protocol description.

Outline of the paper. Section 2 introduces a dialect of the spi-calculus used for modeling security protocols. Section 3 introduces causal graphs. Section 4 defines the abstract interpretation framework and states the soundness results. Section 5 presents the safety results. Section 6 discusses further related work and Section 7 concludes.

2 ρ -spi Calculus

The ρ -spi calculus [11, 12] derives from the spi calculus [4] and inherits many of the features of *Lysa* [8], a dialect of the spi calculus specifically tailored to the analysis of authentication protocols. The ρ -spi calculus differs from both calculi in several respects: it associates principal identities to processes; it syntactically binds keys to their owners; and it provides new authentication-specific constructs. In this paper, we consider a novel dialect of ρ -spi in which encryptions and decryptions are performed on-the-fly when sending and receiving messages, respectively. This dialect in particular links protocol specifications more tightly to their

Table 1 (Our Dialect of) the ρ -spi Calculus

Name	Process
$a ::= I, J, A, B, E$	$P, Q ::= \text{new}(n).P$
$ n, m, k_{IJ}$	$ \text{new}^{\mp}(k_I).P$
$ k_I^+$	$ \text{in}(M).P$
$ k_I^-$	$ \text{out}(M).P$
Term	$ \text{begin}_N^i(A, I, M).P$
$M, N ::= a$	$ \text{end}_N^i(A, I, M).P$
$ x, y, z$	$ A \triangleright P$
$ (M, N)$	$ P Q$
$ \{M\}_u$	$!P$
	$ \mathbf{0}$

Notation: u ranges over names and variables.

informal “graphical” descriptions, which only depict sent and received messages without giving a precise semantics on how messages are parsed and constructed.

2.1 Syntax

The formal syntax of our dialect of the ρ -spi calculus is depicted in Table 1. We presuppose a countable set of *names* partitioned into the set of identities ID , the set of *messages* \mathcal{M} , including the set of shared keys \mathcal{K} , and the set of public and private keys \mathcal{K}^{\mp} . The set ID , ranged over by I and J , is further partitioned into the two sets of *trusted principals* $ID_{\mathcal{P}}$, ranged over by A and B , and *enemies* $ID_{\mathcal{E}}$, ranged over by E . The set \mathcal{K} is composed of keys k_{IJ} shared between I and J . The set \mathcal{K}^{\mp} is partitioned into public and private keys, noted \mathcal{K}^+ and \mathcal{K}^- , respectively. I 's key-pair is composed of a public key k_I^+ and a private key k_I^- , related symbolically by k_I . The set $\mathcal{K}_{\mathcal{E}}$ contains the keys of malicious parties, i.e., the ones where some $E \in ID_{\mathcal{E}}$ occurs as subscript. For convenience, we syntactically bind keys to their owners, thus assuming that keys are already distributed among protocol participants. Notice that this simplification is not restrictive since, if needed, processes can exchange keys thus modelling session-key distribution. Finally, terms can be paired or encrypted with other terms.¹

Processes (or *protocols*), ranged over by P and Q , behave as follows: $\text{new}(n).P$ generates a fresh name n local to P while $\text{new}^{\mp}(k_I)$ generates a fresh key-pair for I composed of k_I^+ and k_I^- . We presuppose a unique unnamed public channel, the network, from/to which all principals, including intruders, read and send messages. Similarly to *Lysa*, our input primitive may atomically test part of the read message, by employing pattern-matching. If the input term matches the input pattern, then the variables occurring in the pattern are bound to the remaining sub-part of the

¹For the sake of readability, in the rest of the paper we omit brackets: for instance, the nested pair $((a, b), k)$ is written as a, b, k .

term; otherwise the term is not read at all. This mechanism is also used to decrypt received messages on-the-fly and thus constitutes an important novelty compared to the ρ -spi calculus; of course, in order to immediately match a term encrypted with asymmetric cryptography, the correct decryption key has to be specified in the pattern. For giving the intuition of the semantics, which is formally defined in Section 2.2, the process ‘ $\text{in}(n).P$ ’ tries to read a specific name n from the network and, if such a name can be read, no binding occurs and P is executed. This is useful, e.g., to check protocols where nonce n is sent encrypted as challenge and received back in clear as response. As another example, consider ‘ $\text{in}(\{x\}_{k_A^-}).P$ ’. This process reads any ciphertext of the form $\{a\}_{k_A^+}$, decrypts it on the fly, and binds all the free occurrences of x to a in process P . We remark that the key specified in the input pattern is the *decryption* key since for binding x to a the process has to perform a decryption and thus to know the correct decryption key. For easing the presentation, we only consider the instantiation of variables with names: as in [14, 27], we assume that messages are typed so as to distinguish names from the other terms like pairs and ciphertexts.

The $\text{begin}_N^i(A, B, M)$ and $\text{end}_N^i(B, A, M)$ primitives express the *correspondence assertions* [33] in a nonce handshake between A and B . The index i allows for specifying which begin assertion should match a specific end assertion: we assume that the same index is used for at most one begin assertion and one end assertion. The former primitive declares that A is starting a protocol session with B , while the latter declares that B is ending a protocol session in which he believes to have correctly authenticated A : N is the nonce used in the protocol session and M is the authenticated message. Finally, $A \triangleright P$ represents principal A executing process P ; $P|Q$ is the parallel composition of P and Q ; $!P$ indicates an arbitrary number of parallel instances of P , and $\mathbf{0}$ is the null process that does nothing. In the rest of the paper, we will often omit $\mathbf{0}$ from protocol specifications.

Example 1 Let us consider a simple protocol where B encrypts message m , nonce n and his own identifier with A ’s public-key and A acknowledges the reception of the first message by sending back the nonce in clear on the network. The specification in ρ -spi calculus is reported below: for the sake of readability, we depict message exchanges in between the corresponding inputs and outputs.

$$\begin{array}{l}
\text{Resp} \triangleq \\
\text{in}(\{B, x, z\}_{k_A^-}). \\
\text{begin}_x^1(A, B, z). \\
\text{out}(x) \\
\\
\text{Prot} \triangleq \text{new}^\mp(k_A).(B \triangleright !\text{Init} \mid A \triangleright !\text{Resp})
\end{array}
\quad
\begin{array}{l}
\text{Init} \triangleq \\
\text{new}(m).\text{new}(n). \\
\text{out}(\{B, n, m\}_{k_A^+}). \\
\\
\text{in}(n). \\
\text{end}_n^1(B, A, m)
\end{array}$$

The begin assertion says that A confirms the reception from B of message m in a protocol session based on nonce n and, similarly, the end assertion says that B authenticates A receiving m in a protocol session based on n . The goal of this protocol is to guarantee the secrecy of m and the authentication of A with B .

2.2 Operational semantics

Following [9], the ρ -spi calculus comes with a trace-based semantics. Each process primitive has an associated action and we denote with Act the set of all possible actions. The dynamics of the calculus is formalized by means of a transition relation between *configurations*, i.e., pairs $\langle s, P \rangle$, where $s \in Act^*$ is a trace and P is a closed process. In the following, ε denotes the empty trace. Each transition $\langle s, P \rangle \rightarrow \langle s :: t, P' \rangle$ simulates one computation step in P and records the corresponding action t in the trace. We denote by \rightarrow^* a finite sequence of computation steps. In the following, G ranges over ground terms, namely terms containing no variable. Principals do not directly synchronize with each other. Instead, they may receive from the unique channel an arbitrary message known to the environment, which models the Dolev-Yao intruder [17]: the knowledge of the environment associated with a ρ -spi trace, characterized by the judgement $s \vdash G$, is formalized by a set of deduction rules stating that the environment knows all the messages sent on the network, every message which is not restricted in the trace, all public keys, and enemies’ keys. The environment can also construct/destroy pairs and encrypt/decrypt ciphertexts if the appropriate key is known. Formal definitions are postponed to Appendix A.

Definition 1 (Traces) *The set $T(P) \triangleq \{s \mid \exists P' \text{ s.t. } \langle \varepsilon, P \rangle \rightarrow^* \langle s, P' \rangle\}$ is the set of all the traces generated by a finite sequence of transitions from the configuration $\langle \varepsilon, P \rangle$.*

Secrecy and authenticity are defined in terms of ρ -spi traces and processes.

Definition 2 (Secrecy) *A trace s guarantees secrecy of G if and only if $s \not\vdash G$. A process P guarantees secrecy of G if and only if $s \not\vdash G$ for all $s \in T(P)$.*

The weak authenticity property refines the standard *non-injective agreement* property of [28, 33] by making explicit the nonce used in the handshake and the index binding the begin assertion with the corresponding end assertion.

Definition 3 (Weak Authenticity) *A trace s guarantees weak authenticity if and only if $s = s_1 :: \text{end}_{G_1}^i(B, A, G_2) :: s_2$ implies $\text{begin}_{G_1}^i(A, B, G_2) \in s_1$. A process P guarantees weak authenticity if and only if s guarantees weak authenticity for all $s \in T(P)$.*

Intuitively, this guarantees that whenever B authenticates A and the message G_2 in a handshake with the nonce G_1 , then A engaged in a protocol session with nonce G_1 for authenticating G_2 with B . Similarly to the *agreement* property of [28, 33], the notion of strong authenticity requires the freshness of authentication requests.

Definition 4 (Strong Authenticity) *A trace s guarantees strong authenticity if and only if whenever $s = s_1 :: \text{end}_{G_1}^i(B, A, G_2) :: s_2$, we have that $s_1 = s'_1 :: \text{begin}_{G_1}^i(A, B, G_2) :: s'_1$ and $s'_1 :: s''_1 :: s_2$ guarantees strong authenticity. A process P guarantees strong authenticity if and only s guarantees strong authenticity for all $s \in T(P)$.*

2.3 Notational conventions

For easing the presentation of the static analysis technique, we use a number of notational conventions. The message restriction $\text{new}(n).P$ is a binder for the message n , the key-pair restriction $\text{new}^\mp(k_I).P$ is a binder for k_I^+ and k_I^- , namely the key-pair based on k_I , and the input primitive is a binder for the variables occurring in the input term with the exception of the decryption keys. In all cases the scope of the binders is the continuation process. Similarly, $\text{new}(n)$ is a binder for message n and $\text{new}^\mp(k_I)$ is a binder for the key-pair based on k_I and their scope is the continuation trace. The names and variables occurring free and bound in processes and traces are defined as usual. As in companion transition systems, e.g. [10], we implicitly identify processes up to renaming of bound variables and names, i.e., up to α -equivalence. To simplify the definition of the static analysis and following [8], we discipline the α -renaming of bound names. We stipulate that for each message n there is a *canonical representative*, noted $\llbracket n \rrbracket$, and we demand that two messages are α -convertible only when they have the same canonical representative. Similar assumption and notation applies to key-pairs and variables. We write $\llbracket P \rrbracket$ to denote the process obtained by replacing each name and variable in P by the corresponding canonical representative. For convenience, we also require that keys, both symmetric and asymmetric ones, having the same canonical representative depend on the same identifiers, i.e., they have the same subscripts. We assume that the set $\mathcal{M}_{\mathcal{E}}$ of names generated by the environment (cf. rule ENV in Table 7) has a canonical representative which is different from the ones of the names bound in the process. Finally, we assume that the bound names of a process are renamed apart and that they do not clash with the free names; much in the same way variables are assumed to be all distinct. For convenience and without loss of generality, we shall reason on protocol specifications in which every bound name and variable has a different canonical representative.

3 Causal Graphs

A *causal graph* $\mathcal{C} = (\mathcal{N}, \mathcal{E})$ is a finite directed graph with nodes \mathcal{N} and edges \mathcal{E} . In the following, we will often write $\text{nodes}(\mathcal{C})$ to denote the set of nodes in \mathcal{C} and $\text{edges}(\mathcal{C})$ to denote the set of edges in \mathcal{C} . Nodes represent abstractions of process events, while edges track the causality among process events. In the following, we discuss and formalize each of these components.

3.1 Nodes and edges

Nodes, also referred to as abstract processes and ranged over by P and Q , are ρ -spi processes built upon ρ -spi names and the following new name categories:

- *message* \mathcal{E} abstracting messages generated by the environment, identities, public-keys and attackers' keys;
- *labelled names* $a_{(x)}$, which are semantically equivalent to names but allow for tracking variable instantiation. As an example, the instantiation in P of the variable x by the name n yields $P[n_{(x)}/x]$. In fact, labelling is local to sequential processes and does not propagate through concurrent threads.
- *indexed names* $a^{(\text{out}(G).P.i)}$, where $\text{out}(G).P$ is the node outputting a and i the positional index of a in G : they only occur within the environment's knowledge and give a precise characterization of which names have been used by the environment to construct a certain term and, notably, the place (i.e., node and position) in which they are sent on the network.

For distinguishing ρ -spi terms from the ones used in causal graphs, we write the latter by sanserif fonts. In the following, we refer to the set of possibly labelled and indexed names as abstract names. Furthermore, we let G range over ground terms, M, N over terms possibly containing variables, v over abstract names, u over abstract names and variables and G over ground term sets.


Edges connect two nodes, so representing the causality among such nodes. We discern two kinds of causality, namely *intra-thread* and *inter-thread causality*.

Intra-thread causality links nodes within the same thread and is represented by edges of the form $P \rightarrow Q$, linking P to the process Q obtained by reducing P .

Inter-thread causality is due to the synchronization among the process P inputting G and reducing into Q and the nodes $\text{out}(G_i).P_i$ outputting the terms used by the environment to construct G . We call such terms integer components since they are not forged by the environment but simply forwarded. They are crucial

for the soundness of our abstraction, which identifies unbounded protocol sessions into a finite model: the soundness of the security properties proved on top of this abstraction requires to determine when different protocol sessions may be interleaved. This may happen when the environment exploits terms output in different protocol sessions to construct a term that is sent to a principal engaging in another protocol session: in fact, messages within the same integer component prove to abstract messages belonging to the same protocol session, while messages in different integer components may abstract messages belonging to different protocol sessions. This kind of causality is represented by (i) *input edges* of the form $\text{in}(M).P \dashv \boxed{G} \rightarrow Q$, connecting the input process with the process obtained by performing the input via the synchronization point \boxed{G} ; and (ii) *output edges* of the form $\text{out}(G_i) \xrightarrow{G'_i} \boxed{G} \rightarrow Q$, one for each integer component G'_i of the input term. These edges are labelled by the integer component and connect the process outputting such a component to the process obtained by performing the input via the synchronization point \boxed{G} . A synchronization point links an input edge to the output edges required to construct the input term.

Example 2 Let \mathcal{C} be a causal graph containing the nodes $\text{out}(n).Q_1$, $\text{out}(\{m, n\}_{k_A^-}).Q_2$ and $\text{in}(\{x, \{y, z\}_{k_A^+}\}_{k_B^-}).Q_3$. Let us define the following indexes: $i_1 = (\text{out}(n).Q_1, 1)$, $i_2 = (\text{out}(\{m, n\}_{k_A^-}).Q_2, 1)$, $i_3 = (\text{out}(\{m, n\}_{k_A^-}).Q_2, 2)$, and $i_4 = (\text{out}(\{m, n\}_{k_A^-}).Q_2, 3)$. The environment may combine the two output terms into $\{n^{i_1}, \{m^{i_2}, n^{i_3}\}_{k_A^- i_4}\}_{\mathcal{E}}$, which matches the input pattern $\{x, \{y, z\}_{k_A^+}\}_{k_B^-}$ with substitution $\sigma = [n_{(x)}/x, m_{(y)}/y, n_{(z)}/z]$: the substitution σ is used to instantiate the free variables in the process Q_3 following the input pattern. Thus \mathcal{C} contains the following nodes and edges, where $G = \{n_{(x)}^{i_1}, \{m_{(y)}^{i_2}, n_{(z)}^{i_3}\}_{k_A^- i_4}\}_{\mathcal{E}}$:



Notice that the synchronization point tracks both the positional indexes of the names occurring in the input term and the variable assignment induced by the input pattern.

3.2 Causal graphs as abstraction of ρ -spi processes

Before formalizing the causal graph generation, we introduce some useful notations. We write $\lfloor M \rfloor$, $\lceil M \rceil$ and $\llbracket M \rrbracket$

to denote the term obtained from M by label erasure, index erasure, and both label and index erasure, respectively.

Function $\text{outmsg}(\text{out}(G).P)$ yields the term obtained from G by erasing the labels in G and indexing each name a in G with the pair composed of process $\text{out}(G).P$ and the positional index of a in G . This function is naturally extended to node sets, thus characterizing the set of output terms: more formally, $\text{outmsg}(\mathcal{N}) = \{\text{outmsg}(\text{out}(G).P) \mid \text{out}(G).P \in \mathcal{N}\}$. Notice that labels are local and do not propagate into the environment's knowledge.

The knowledge of the abstract environment is formalized by the judgement $\mathcal{G} \vdash G$ (cf. Table 8 of Appendix B), meaning that the environment can construct G given the knowledge of the terms in \mathcal{G} . The environment knows every term in \mathcal{G} and the special name \mathcal{E} , it can construct and destruct pairs and encrypt and decrypt terms provided that it knows the encryption and decryption keys, respectively. For reducing the number of terms known to the environment, we prevent the environment from deriving identities, public keys and enemies' keys, which abstract the same ρ -spi terms as \mathcal{E} . Function $\text{bind} : (M, G) \mapsto (G', \sigma)$, reported in Table 8 of Appendix B, defines the pattern-matching among terms. This function takes as input a term M and a ground term G and, if the two terms match, yields the term G' obtained by labelling G according to the variables in M and the substitution σ expressing the variable instantiation induced by the pattern-matching. If pattern-matching fails, bind returns \uparrow . For example, $\text{bind}(x, n^5) = (n_{(x)}^5, [n_{(x)}/x])$. Function thread , defined below, yields the set of threads in P , abstracting away from identifiers and replications.

$$\text{thread}(P) \triangleq \begin{cases} \text{thread}(P_1) \cup \text{thread}(P_2) & \text{if } P = P_1 | P_2 \\ \text{thread}(Q) & \text{if } P \in \{A \triangleright Q, !Q\} \\ \{P\} & \text{otherwise} \end{cases}$$

Function $\text{int}(\mathcal{N}, G)$ yields the set of integer terms in G , that is the largest subterms of G that have not been generated by the environment, each of them paired with the outputting process. In fact, integer terms are either names or ciphertexts whose encryption key is unknown to the environment.

$$\text{int}(\mathcal{N}, G) \triangleq \begin{cases} \{(v^{(P,j)}, P)\} & \text{if } G = v^{(P,j)} \\ \emptyset & \text{if } G = \mathcal{E} \\ \text{int}(\mathcal{N}, G_1) \cup \text{int}(\mathcal{N}, G_2) & \text{if } G = (G_1, G_2) \\ \{(G, P)\} & \text{if } G = \{G'\}_{v^{(P,j)}} \wedge \text{outmsg}(\mathcal{N}) \not\vdash v^{(P,j)} \\ \text{int}(\mathcal{N}, G') \cup \text{int}(\mathcal{N}, v') & \text{if } G = \{G'\}_{v'} \wedge \text{outmsg}(\mathcal{N}) \vdash v' \end{cases}$$

The next definition characterizes the causal graph associated with an abstract process. Here and throughout this paper, we write $P =_\alpha Q$ to say that P and Q are α -equivalent and $\llbracket P \rrbracket$ to denote the process obtained by replacing each name and variable in P with the corresponding canonical representative.

Definition 5 (Causal Graph) *The causal graph $(\mathcal{N}, \mathcal{E})$ associated with P , written as $(\mathcal{N}, \mathcal{E}) = \text{graph}(P)$, is given by the least \mathcal{N}, \mathcal{E} satisfying the following conditions:*

- (i) $\text{thread}(P) \subseteq \mathcal{N}$
- (ii) $p.P \in \mathcal{N} \wedge Q \in \text{thread}(P) \wedge p \neq \text{in}(\cdot) \Rightarrow Q \in \mathcal{N} \wedge p.P \rightarrow Q \in \mathcal{E}$
- (iii) $\text{in}(M).P \in \mathcal{N} \wedge \text{outmsg}(\mathcal{N}) \vdash G \wedge \text{bind}(M, G) = (G', \sigma) \wedge P' \in \text{thread}(P\sigma) \wedge (G', \text{out}(G_1).P_1) \in \text{int}(\mathcal{N}, G') \Rightarrow \exists Q \in \mathcal{N} \text{ s.t. } Q =_{\alpha} P' \wedge \text{in}(M).P \xrightarrow{\boxed{G'}} Q \in \mathcal{E} \wedge \text{out}(G_1).P_1 \xrightarrow{G_1} \boxed{G'} \rightarrow Q \in \mathcal{E}$
- (iv) $\forall \{\text{new}(n).P, \text{new}(m).Q\} \subseteq \mathcal{N}, \llbracket \text{new}(n).P \rrbracket = \llbracket \text{new}(m).Q \rrbracket \Leftrightarrow n = m$

Function $\text{graph} : P \mapsto (\mathcal{N}, \mathcal{E})$ is in fact a closure operator on causal graphs ordered by inclusion. In particular, (i) the threads in P are part of the set \mathcal{N} of nodes; (ii) if $p.P$ is a node in the causal graph, with $p \neq \text{in}(\cdot)$, then the processes in $\text{thread}(P)$, which are connected to $p.P$ by direct edges, belong to the set of nodes; and (iii) if $\text{in}(M).P$ is a node and the environment knows a term G matching M with substitution σ , then the set of nodes contains the processes in $\text{thread}(P\sigma)$, up to α -equivalence, and the set of edges contains the input edges connecting $\text{in}(M).P$ to the processes in $\text{thread}(P\sigma)$ and the output edges linking the output processes used by the environment to construct G to the processes in $\text{thread}(P\sigma)$. Intuitively, the input of a term is preceded by the output of *all* messages required by the environment to construct such a term. This point will be clarified later on, when formalizing the causality relation expressed by causal graphs. Condition (iv) rules the α -renaming of bound names possibly introduced in the causal graph by condition (iii). This α -renaming allows the abstract domain to distinguish the names generated in different threads, thus making the analysis more precise. Unfortunately, this may introduce an infinite number of names in the knowledge of the environment and possibly an infinite number of nodes in the causal graph. We tackle this problem by requiring that whenever two restriction nodes occur in the causal graph and the restricted names differ because of α -renaming, then at least one variable is instantiated with names belonging to different equivalence classes in the two continuation processes. The idea is to apply α -renaming of bound names in the processes following an input for every different input term, where the difference among input terms is deemed up to name equivalence classes. This enhances the precision of the analysis, yet guaranteeing the finiteness of the causal graph.

We say that two graphs are isomorphic when they are obtained by α -renaming of bound names. As stated by the following proposition, every abstract process admits a unique causal graph up to isomorphism. In the following,

we write $\text{length}(P)$ to denote the number of primitives in P and $\text{length}(\mathcal{N})$ to denote the maximal length of the processes in \mathcal{N} , namely the number n such that $\text{length}(Q) = n$, for some $Q \in \mathcal{N}$, and, for every $Q' \in \mathcal{N}$, $\text{length}(Q') \leq n$.

Proposition 1 (Uniqueness) *For every process P , there exists a unique graph (P) up to isomorphism.*

Proof. Consider the function $\Phi_P(\mathcal{N}, \mathcal{E})$ yielding the least $\mathcal{N}', \mathcal{E}'$ satisfying the three conditions of Definition 5. For proving the thesis, we prove that, for every \mathcal{N} and \mathcal{E} , there exists a unique least fixpoint of $\Phi_P(\mathcal{N}, \mathcal{E})$: this trivially implies the thesis. Notice that Φ_P is monotonous over node and edge sets ordered by inclusion.

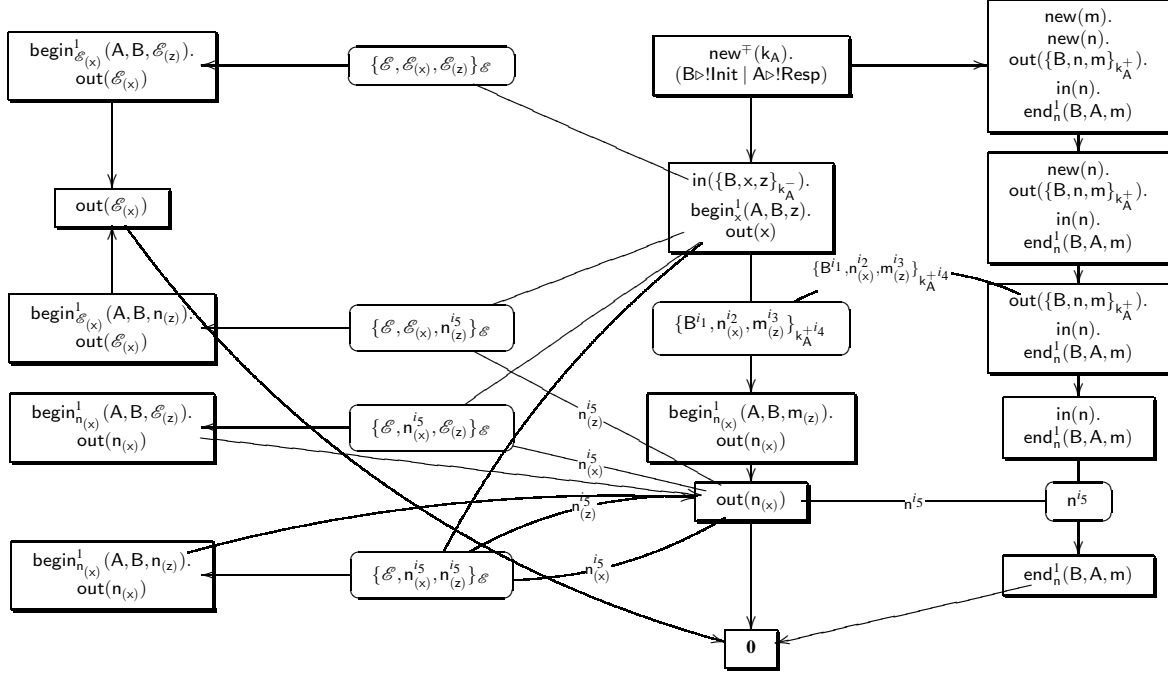
It is easy to see that if $\Phi_P(\mathcal{N}, \mathcal{E}) = \mathcal{N}', \mathcal{E}'$, then $\text{length}(\mathcal{N}) = \text{length}(\mathcal{N}')$. The set of variables occurring in the processes of \mathcal{N}' is contained into the set of variables occurring in the processes of \mathcal{N} . The situation for names is different, since some variable might be instantiated with \mathcal{E} and, notably, new names might be introduced by α -renaming (condition (iii)). However, because of condition (iv), the number of α -renamed copies of each name is bound by N^X , where N and X are the number of restrictions and variables in \mathcal{N} , respectively. The set of processes of finite length and composed of a finite set of names and variables is finite, as well as the set of edges. By Knaster-Tarski theorem, Φ_P has a unique least fixpoint. \square

The following corollary says that the size of the graph grows exponentially with the protocol *specification*, which is however fixed in advance, regardless of the number of considered sessions and the protocol run-time behavior. We remark that this result refers to the worst case: in practice, the number of nodes for the protocols we have considered so far does not exceed 100 nodes and the analysis always terminated within a few seconds.

Corollary 1 (Size of causal graphs) *Let \mathcal{C} be a causal net and P a process such that $\mathcal{C} = \text{graph}(P)$. Let N and X be the number of restrictions and variables in P , respectively. Then $|\text{nodes}(\mathcal{C})| = O(\text{length}(P) * N^{X^2})$.*

Example 3 The causal graph associated with protocol *Prot* of Example 1 is depicted in Table 2. The rounded boxes represent synchronization points while the other ones abstract processes. The analysis of the causal graph gives us some interesting information about the run-time behavior of protocol participants. Even if the protocol is simple, it turns out to be interesting since the attackers know both the public-key used for encrypting the first message and, after the second message, the nonce used in the protocol session. This increases the number of actions at their disposal and, consequently, the number of nodes and edges in the causal graph. In general, fixed the number of message exchanges, the causal graph for protocols preserving the secrecy of

Table 2 Causal graph $graph(Prot)$



Legend: for every $j \in [1, 4]$, $i_j = (out(\{B, n, m\}_{k_A}^+).in(n).end_n^1(B, A, m), j)$; $i_5 = (out(n_{(x)}), 1)$

messages and relying on digital signature or symmetric-key cryptography is typically simpler than the one of protocols based on public-key cryptography and possibly exposing some messages to the attackers. Let \mathcal{N} be the set of nodes in the causal graph: we can see that $outmsg(\mathcal{N}) \not\subseteq m^i$ for any index i and this intuitively means that the authenticated message is kept secret, as expected. Furthermore, the only process asserting an end event is $end_n^1(B, A, m)$. This is preceded by the output of $n_{(x)}$ (inter-thread causality), which is in turn preceded (intra-thread causality) by node $begin_{n_{(x)}}^1(A, B, m(z)).out(n_{(x)})$. Note that the output of $n_{(x)}$ may be also preceded (intra-thread causality) by $begin_{n_{(x)}}^1(A, B, E(z)).out(n_{(x)})$ and $begin_{n_{(x)}}^1(A, B, n(z)).out(n_{(x)})$. These events are enabled by the environment forging the messages $\{E, n^{i_5}, E\}_E$ and $\{E, n^{i_5}, n^{i_5}\}_E$, respectively (cf. the corresponding synchronization points), and sending them to the responder. This may happen only if the environment knows n and these processes are thus preceded (inter-thread causality) by the output of $n_{(x)}$. This cycle in the graph tells us that the environment can actively interact with the responder only after the responder has received the message generated by the initiator and “asserted” $begin_{n_{(x)}}^1(A, B, m(z))$, according to the intended protocol run. Intuitively, this means that the causal graph guarantees authenticity. Finally, the portion of the graph in the upper-left corner shows that the environment can forge mes-

sages in which the message E replaces the nonce, but the resulting processes do not increase the environment’s knowledge as they eventually output E , which is already known to the environment.

3.3 Paths

Edges describe the causality among process events and each edge may be naturally associated with an action: for instance, if a causal graph contains the edge $out(n).P \rightarrow P$, then the process P is causally preceded by the reduction of the process $out(n).P$ with action $out(n)$. In general, we express the causality dependency among process events through *paths*, which are sequences of *abstract actions* expressing process events. Abstract actions extend p-spi ones with the new action $in_{com}(G_1, G_2)$, which is used for tracking inputs involved in inter-thread causality: G_1 is the integer component and G_2 is the input pattern, in which variables are replaced with the actual received messages.

We now argue on the number of paths associated with a node. We have mentioned that the input of a term is causally preceded by the outputs of *all* the integer components needed by the environment to construct such a term. Thus an input is preceded by a set of paths. However, the same process may be generated by the input of different terms, say G_1 and G_2 . Then such a process is preceded by either all the outputs needed by the environment to con-

struct G_1 or all the outputs needed to construct G_2 . This is the reason why the paths associated with a node are actually a set of path sets, meaning that the node is causally preceded by *all* the paths in *one* of its path sets. Notice that, as result of our abstraction which essentially collapses an unbounded number of instances of each principal into a strand of nodes related by intra-thread causality, causal graphs may contain cycles. For this reason, when evaluating the paths preceding a node, we need to avoid loops: this is achieved by traversing input edges only once, thus abstracting away from cycles in the causal graph which, in fact, do not alter the causality among nodes. In the following, we write $\{s_1, \dots, s_n\} :: \tau$ to denote $\{s_1 :: \tau, \dots, s_n :: \tau\}$ and $out_{\text{edg}}(G, Q, \mathcal{C})$ to denote the set of output edges in \mathcal{C} incoming in Q via the synchronization point \boxed{G} . Function $paths : \mathcal{C}, P, \mathcal{E} \mapsto \mathcal{S}$, defined below, yields the set \mathcal{S} of paths sets that are associated with the node P and do not traverse the edges in \mathcal{E} . We often write $paths(\mathcal{C}, P, \emptyset)$.

Definition 6 (Paths) Let \mathcal{C} be a causal graph and $Q \in nodes(\mathcal{C})$. The paths preceding Q in \mathcal{C} and not traversing the edges in \mathcal{E} , written $paths(\mathcal{C}, Q, \mathcal{E})$, are defined as the least \mathcal{S} such that

- (a) Q has no incoming edge $\Rightarrow \mathcal{S} = \{\emptyset\}$
- (b) $p.P \rightarrow Q \in edges(\mathcal{C}) \wedge S \in paths(\mathcal{C}, p.P, \mathcal{E}) \Rightarrow S :: p \in \mathcal{S}$
- (c) $in(M).P \xrightarrow{\boxed{G}} Q \in edges(\mathcal{C}) \setminus \mathcal{E} \wedge bind(M, [G]) = (G, \sigma) \wedge \bigcup_{i \in [1, n]} \{out(G_i).P_i \xrightarrow{G_i} \boxed{G} \rightarrow Q\} = out_{\text{edg}}(G, Q, \mathcal{C}) \wedge S \in paths(\mathcal{C}, in(M).P, \mathcal{E} \cup \{in(M).P \xrightarrow{\boxed{G}} Q\}) \wedge S_i \in paths(\mathcal{C}, out(G_i).P_i, \mathcal{E} \cup \{in(M).P \xrightarrow{\boxed{G}} Q\}) \Rightarrow S :: in(M\sigma) \cup \bigcup_{i \in [1, n]} S_i :: out(G_i) :: in_{\text{com}}(G'_i, M\sigma) \in \mathcal{S}$

If the node Q is preceded by intra-thread causality by $p.P$ (condition (a)), then the path sets associated with Q are the ones associated with $p.P$, each of them extended with the action p . If Q is obtained by the reduction of process $in(M).P$ receiving G and the input edge $in(M).P \xrightarrow{\boxed{G}} Q$ is not in \mathcal{E} , then the paths associated with Q are the ones associated with $in(M).P$, each of them extended with the action $in(M\sigma)$, plus the ones associated with the processes outputting the terms used to construct G (edges $out(G_i).P_i \xrightarrow{G_i} \boxed{G} \rightarrow Q$), extended with $out(G_i) :: in_{\text{com}}(G'_i, M\sigma)$. Notice that the edge $in(M).P \xrightarrow{\boxed{G}} Q$ is inspected only once, thus avoiding loops due to cycles in the graph.

3.4 Names and sessions

Since a causal graph abstracts an unbounded number of protocol sessions, an interesting issue related to our abstrac-

tion is that different occurrences in a path of the same name n might actually abstract different ρ -spi names, one for each abstracted protocol session. The problem is that the environment may exploit messages generated in different protocol sessions so as to forge a message which is then used to interact with another session: in fact, this kind of interleaving may break the security goals of the protocol. According to the definition of causal graphs, the labels of the output edges express the integer components of the input term. This information can be used for soundly characterizing which names abstract over the same protocol session and, in particular, which names abstract the same ρ -spi name. Intuitively, it is sufficient to check for every pair of output and input actions related by inter-thread causality which messages belong to the integer component and which do not: the former abstract over the same protocol session while the latter may belong to different protocol sessions. Here and throughout this paper, we write $msgs(M)$ to denote the set of messages in M , including those labelled. We also write $\pi^i(G)$ to denote the name occurring in the i -th position of G and $\pi^i(s)$ to denote the i -th action in s . Finally, $|s|$ denotes the number of actions in s .

Table 3 introduces the deduction system for judgement $s \vdash (v, i) = (v', j)$, meaning that names v and v' , occurring in the i -th and j -th action of s respectively, abstract over the same ρ -spi name. Intuitively, the concretization of a path yields the set of ρ -spi traces obtained by instantiating the abstract names occurring therein: if $s \vdash (v, i) = (v', j)$, then the occurrences of v and v' in the i -th and j -th action of s , respectively, are instantiated with the same ρ -spi name (see Section 4 for more detail). Rule AX says that the occurrences of the same message within a single action abstract the same ρ -spi name. Rule INTRA says that the occurrences of a message within actions related by intra-thread causality abstract the same ρ -spi name, since such actions belong to the same thread and thus abstract over the same protocol session. Rule INTER-BIND says that messages received within an integer component, that is messages not manipulated by the attacker, abstract over the same protocol session. Rule INTER-MATCH says that the pattern-matching of name n allows for recovering a protocol session described in a previous part of the path and, more precisely, it allows to inherit the equality constraints holding at the time of n 's restriction. Finally, TRANS, SYM, and PREFIX make the relation transitive, symmetric, and closed by prefixes, respectively.

Example 4 Let us consider the following path, which represents the intended protocol behavior:

$$s \triangleq new^{\mp}(k_A) :: new(m) :: new(n) :: out(\{B, n, m\}_{k_A^+}) :: in_{\text{com}}(\{B^{i_1}, n_{(x)}^{i_2}, m_{(z)}^{i_3}\}_{k_A^{+i_4}}, \{B, n_{(x)}, m_{(z)}\}_{k_A^-}) :: begin_{n_{(x)}}^1(A, B, m_{(z)}) :: out(n_{(x)}) :: in_{\text{com}}(n^{i_5}, n) :: end_n^i(B, A, m)$$

Table 3 Equality among abstract terms

$\frac{\text{AX}}{s :: t \mid = i \quad v \in \text{msgs}(t)} s :: t \vdash (v, i) = (v, i)$	
$\frac{\text{INTRA}}{ s = i \quad t \neq \text{incom}(G, G') \quad s \vdash (v, i) = (v, j)} s :: t \vdash (v, i + 1) = (v, i)$	
$\frac{\text{INTER-BIND}}{v^{(P,j)} \in \text{msgs}(G_2) \quad \pi^j(G_1) = v' \quad s = i} s :: \text{out}(G_1) :: \text{incom}(G_2, G) \vdash (v, i + 2) = (v', i + 1)$	
$\frac{\text{INTER-MATCH}}{n \in \text{msgs}(G) \quad s = s_1 :: \text{new}(n) :: s_2 \quad s = i \quad s_1 = j \quad s \vdash (n, i + 2) = (n, j + 1) \quad s_1 \vdash (v, j) = (v', j')}{s :: \text{out}(G_1) :: \text{incom}(G_2, G) \vdash (v, i + 2) = (v, j)}$	
$\frac{\text{TRANS}}{s \vdash (v, i) = (v', i') \quad s \vdash (v', i') = (v'', i'')}{s \vdash (v, i) = (v'', i'')}$	
$\frac{\text{SYM}}{s \vdash (v, i) = (v', i')}{s \vdash (v', i') = (v, i)}$	$\frac{\text{PREFIX}}{s \vdash (v, i) = (v', i')}{s :: t \vdash (v, i) = (v', i')}$

The messages occurring in the begin and end assertions abstract over the same protocol session. In particular, we can prove $s \vdash (n, 9) = (n_{(x)}, 6)$ via INTRA and INTER-BIND. Similarly, we can prove $s \vdash (m, 9) = (m, 2)$ via INTRA and INTER-MATCH, and $s \vdash (m_{(z)}, 6) = (m, 2)$ via INTRA, INTER-BIND and PREFIX. Finally, by SYMM and TRANS we get $s \vdash (m, 9) = (m_{(z)}, 6)$.

In Section 3.3, we have defined the paths associated to a node in a causal graph and, for avoiding loops due to cycles in the graph, we have requested that input edges are inspected only once. This approximation does not affect the causality among nodes but it might affect the equality constraints. For this reason, we only consider a class of causal graphs, called *cycle-invariant* graphs, for which cycles do not affect the equality constraints of the involved paths.

Definition 7 (Cycle-invariance) *We say that \mathcal{C} is cycle-invariant if and only if for every $\{\text{in}(M).P \mid \boxed{G} \rightarrow Q, \text{out}(G_1).P' \xrightarrow{G_2} \boxed{G} \rightarrow Q\} \subseteq \text{edges}(\mathcal{C}), S_{out} \in \text{paths}(\mathcal{C}, \text{out}(G_1).P'), S_{in} \in \text{paths}(\mathcal{C}, \text{in}(M).P)$ there exists $S \in \text{paths}(\mathcal{C}, Q)$ s. t. the following conditions hold:*

- for every $s_1 = s' :: \text{in}(G') \in \mathcal{S}$ there exists $s_2 \in S_{in}$ such that either $s_2 = s'$ or $s_2 = s' :: \text{in}(G') :: s'' \in S_{in}$ and $s_1 \vdash (v, j) = (v', j') \Rightarrow s_2 \vdash (v, i) = (v', j')$, with $|s_1| = j$ and $|s_2| = i$.

- for every $s_1 = s' :: \text{out}(G_1) :: \text{incom}(G_2, G') \in \mathcal{S}$ there exists $s_2 \in S_{out}$ such that either $s_2 = s'$ or $s_2 = s' :: \text{out}(G_1) :: \text{incom}(G_2, G') :: s''$ and $s_1 \vdash (v, j) = (v', j') \Rightarrow s_2 \vdash (v, i) = (v', j')$, with $|s_1| = j$ and $|s_2| = i$.

For example, *graph(Prot)* is cycle-invariant in that input nodes do not belong to cycles, the only output node within cycles is $\text{out}(n_{(x)})$, and cycles in the causal graph do not affect the equality constraints on $n_{(x)}$.

4 Abstract Interpretation

In this section we illustrate the relation between causal graphs and ρ -spi semantics. This is formalized by a concretization function, defined on abstract terms, paths, path sets and causal graphs.

4.1 Concretization of causal graphs

Terms The relation between ρ -spi terms and abstract terms is formally defined in Table 4 by the concretization function $\gamma_{\text{trm}} : M \mapsto \{M_1, \dots, M_n\}$. Abstract identities are instantiated with the corresponding ρ -spi identity. The concretization of an abstract message yields the set of messages having the same canonical representative. Similar reasoning applies to variables, public and private keys. The special name \mathcal{E} abstracts over identities, public keys, messages possibly generated by the environment and attackers' keys. Finally, the concretization of the remaining terms is given by instantiating the names and the variables occurring therein.

Paths The concretization function $\gamma_{\text{path}} : s \mapsto \{(s_1, \sigma_1), \dots, (s_n, \sigma_n)\}$ takes as input a path and yields a set of pairs composed of a ρ -spi trace s_i and a substitution σ_i from abstract names to ρ -spi names. More precisely, (i) each trace is obtained by instantiating the names occurring in each action of the path so as to satisfy the equality constraints associated to the path and (ii) the substitution tracks the instantiation of the abstract messages occurring in the actions of the last thread, namely the largest suffix of s related by intra-thread causality. We recall that a path may “cross” different threads and, as a matter of fact, $s \vdash (v, k) = (v', k')$, with k being the length of s , only if v is an abstract name occurring in the last thread.

Example 5 One of the possible concretizations of the path s described in Example 4 is reported below:

$$\begin{aligned} & \text{new}^{\mp}(k_A) :: \text{new}(m) :: \text{new}(n) :: \text{out}(\{B, n, m\}_{k_A^+}) :: \\ & \text{in}(\{B, n, m\}_{k_A^-}) :: \text{begin}_n^1(A, B, m) :: \text{out}(n) :: \\ & \text{in}(n) :: \text{end}_n^1(B, A, m) \\ & \sigma : n \mapsto n, m \mapsto m \end{aligned}$$

Table 4 Concretization functions

$$\begin{aligned}
\gamma_{\text{trm}}(M) &\triangleq \begin{cases} \{I\} & \text{if } M = I \\ \{n \mid \llbracket n \rrbracket = \llbracket m \rrbracket\} & \text{if } M = m \\ \{y \mid \llbracket y \rrbracket = \llbracket x \rrbracket\} & \text{if } M = x \\ \{k'_j{}^+ \mid \llbracket k'_j \rrbracket = \llbracket k_I \rrbracket\} & \text{if } M = k_I^+ \\ \{k'_j{}^- \mid \llbracket k'_j \rrbracket = \llbracket k_I \rrbracket\} & \text{if } M = k_I^- \\ \mathcal{ID} \cup \mathcal{K}^+ \cup \mathcal{M}_E \cup \mathcal{K}_E & \text{if } M = \mathcal{E} \\ \gamma_{\text{trm}}(a) & \text{if } M \in \{a_{(x)}, a^i, a^i_{(x)}\} \\ \{M\sigma \mid \forall u \in \text{dom}(\sigma), \sigma(u) \in \gamma_{\text{trm}}(u)\} & \text{otherwise} \end{cases} \\
\gamma_{\text{path}}(s) &\triangleq \{(s, \sigma) \mid |s| = |s| = n \wedge \begin{array}{l} (i) \quad \forall i \in [1, n], \exists \sigma_i \text{ s.t. } \pi^i(\bar{s})\sigma_i = \pi^i(s) \wedge (s \vdash (v, j) = (v', j') \Rightarrow \sigma_j(v) = \sigma_{j'}(v')) \wedge \\ (ii) \quad (s \vdash (v, n) = (v', k) \wedge \sigma_k(v') = a) \Leftrightarrow \sigma(v) = a \end{array}\} \\
\gamma_{\text{pset}}(S) &\triangleq \{(\{s_1, \dots, s_n\}, \sigma) \mid S = \{s_1, \dots, s_n\} \wedge \forall i \in [1, n], (s_i, \sigma_i) \in \gamma_{\text{path}}(s_i) \wedge \sigma = \bigsqcup_{i=1, n} \sigma_i \neq \uparrow\} \\
\gamma_{\text{net}}(\mathcal{C}) &\triangleq \{(s, P) \mid \begin{array}{l} (i) \quad \forall i \in [1, n], \exists Q_i \in \text{nodes}(\mathcal{C}), S_i \in \text{paths}(\mathcal{C}, Q_i), (S_i, \sigma_i) \in \gamma_{\text{pset}}(S_i) \text{ s.t. } S_i = S_j :: \pi^i(s), \text{ with } j < i \wedge \\ (ii) \quad \forall Q \in \text{thread}(P), \exists i \in [1, n], \sigma \text{ s.t. } Q = Q_i\sigma \wedge \sigma \neq \uparrow \end{array}\}
\end{aligned}$$

Notation: $\pi^i(s)$ yields the i -th action in s

\bar{s} yields the trace obtained from s by replacing each occurrence of $\text{in}_{\text{com}}(G_1, G_2)$ with $\text{in}(G_2)$

$\sigma_1 \sqcup \sigma_2 = \sigma_1 \cup \sigma_2$ if $\sigma_1 \neq \uparrow \wedge \sigma_2 \neq \uparrow \wedge u \in \text{dom}(\sigma_1) \cap \text{dom}(\sigma_2) \Rightarrow \sigma_1(u) = \sigma_2(u)$

$\sigma_1 \sqcup \sigma_2 = \uparrow$ otherwise

Since $s \vdash (n_{(x)}, 6) = (n, 9)$, $n_{(x)}$ and n are instantiated with the same name n . Similarly, since $s \vdash (m_{(z)}, 6) = (m, 9)$, m and $m_{(z)}$ are instantiated with the same name m . Notice also that the path guarantees authenticity and the trace guarantees strong authenticity.

Path sets The concretization of a path set is given by the concretization of the paths occurring therein and by the union of the corresponding instantiations. It is worth to mention that the instantiation of the messages in last thread has to be the same ($\bigsqcup_{i=1, n} \sigma_i \neq \uparrow$): the union $\sigma_1 \sqcup \sigma_2$ succeeds only when the abstract names substituted by both σ_1 and σ_2 are bound to the same ρ -spi name.

Configurations The concretization function $\gamma_{\text{net}} : \mathcal{C} \mapsto \{\langle s_1, P_1 \rangle, \dots, \langle s_n, P_n \rangle\}$ takes as input a causal graph and yields a set of ρ -spi configurations satisfying two constraints, the former concerning traces and the latter concerning processes. In the following, we let S range over ρ -spi trace sets.

i for every action τ_i in the trace there exists a node Q_i in the causal graph, a path set S_i associated to Q_i and a trace set S_i such that S_i is an instantiation of S_i ($(S_i, \sigma_i) \in \gamma_{\text{pset}}(S_i)$), where S_i is a trace set associated to a preceding action in the trace, extended with τ_i .

ii for every thread Q of the process in the ρ -spi configuration, there exists a node Q_i in the causal graph, a path set S_i , a trace set S_i and an instantiation σ_i such

that $Q = Q_i\sigma$, where σ is compatible with the name instantiation σ_i given by the concretization of the path set associated to τ_i in the last thread.

Intuitively, condition (i) requires that the actions in the trace respect the causality paths induced by the causal graph and condition (ii) requires that every process has an abstraction in the causal graph and the trace respects the causality paths associated to such a node.

Finally, we state the soundness results of the abstract interpretation. The following theorem says that the causal graph generated from a process is an abstraction of the configuration composed of such a process and the empty trace.

Theorem 1 (Soundness) *If $\text{graph}(P) = \mathcal{C}$, then $\langle \varepsilon, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$.*

The following theorem says that the set of configurations abstracted by a causal graph is closed under process reduction, i.e., causal graphs are a sound abstraction of the ρ -spi semantics.

Theorem 2 (Preservation) *Let \mathcal{C} be cycle-invariant. If $\langle s, P \rangle \rightarrow \langle s', P' \rangle$ and $\langle s, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, then $\langle s', P' \rangle \in \gamma_{\text{net}}(\mathcal{C})$.*

5 Safety Results

In this section we state the safety results of our static analysis technique. As stated by the following definition, a causal graph guarantees the secrecy of a name v , if the abstract environment cannot deduce any term which is equal to v up to index erasure.

Definition 8 (Abstract secrecy) A causal graph \mathcal{C} guarantees the secrecy of v if and only if $\text{outmsg}(\text{nodes}(\mathcal{C})) \vdash v'$ implies $\lceil v \rceil \neq \lceil v' \rceil$.

As formalized below, a path guarantees authenticity if every end assertion is preceded by a corresponding begin assertion and the terms occurring therein are related by equality constraints. This guarantees that authenticity carries over the action sequences abstracted by the path. A causal graph guarantees authenticity if for every $\text{end}_{G_1}^i(A, B, G_2).P$ and $S \in \text{paths}(\mathcal{C}, \text{end}_{G_1}^i(A, B, G_2).P)$, there exists a path in S containing a suitable begin assertion. In the following, we write $G \simeq_{s,i,j} G'$, read as G is equivalent to G' in s , if the two terms are equal component-wise, i.e., for every k such that $\pi^k(G) = v$ and $\pi^k(G') = v'$, we have that v and v' are either the same identity or possibly labelled messages satisfying $s \vdash (v, i) = (v', j)$. This guarantees that the two terms abstract the same ρ -spi term.

Definition 9 (Abstract Authenticity) A path s guarantees authenticity if and only if for every j such that $\pi^j(s) = \text{end}_{G_1}^i(A, B, G_2)$, there exist $G'_1, G'_2, j' < j$ s.t. $\pi^{j'}(s) = \text{begin}_{G'_1}^i(B, A, G'_2)$, $G_1 \simeq_{s,j,j'} G'_1$ and $G_2 \simeq_{s,j,j'} G'_2$.

A causal graph \mathcal{C} guarantees authenticity iff for every $\text{end}_{G_1}^i(A, B, G_2).P \in \text{nodes}(\mathcal{C})$ and $S \in \text{paths}(\mathcal{C}, \text{end}_{G_1}^i(A, B, G_2).P)$, there exists $s \in S$ s.t. $s :: \text{end}_{G_1}^i(A, B, G_2)$ guarantees authenticity.

For example, the path d in Example 4 guarantees authenticity as the end assertion is preceded by a corresponding begin assertion and the messages occurring therein are equivalent. As a matter of fact, every $S \in \text{paths}(\text{graph}(\text{Prot}), \text{end}_n^1(B, A, m))$ contains a path guaranteeing authenticity. Thus $\text{graph}(\text{Prot})$ guarantees authenticity. The following theorems state that causal graphs constitute a sound model for the static verification of secrecy and authenticity. In particular, the next theorem says that if the causal graph associated with P guarantees the secrecy of v , then P guarantees the secrecy of any concretization of v .

Theorem 3 (Secrecy) Let P be a process and $\mathcal{C} = \text{graph}(P)$ a cycle-invariant causal graph. If \mathcal{C} guarantees the secrecy of v then P guarantees the secrecy of any $v \in \gamma_{\text{trm}}(v)$.

For instance, the causal graph of Example 3 guarantees the secrecy of m and, consequently, the protocol of Example 1 guarantees the secrecy of the authenticated message in every protocol session. In the following we say that a process is *nonce linear* iff (i) every end assertion has the form $\text{end}_n^i(I, J, M)$ and is preceded by $\text{new}(n)$ and (ii) if $\text{end}_n^i(I, J, M)$ occurs inside the scope of a replication then $\text{new}(n)$ occurs inside the scope of the same replication. In fact, this syntactic condition suffices to prove that authenticity on causal graphs implies strong authenticity on the ρ -spi processes abstracted by such causal graphs.

Theorem 4 (Authenticity) Let P be a process and $\mathcal{C} = \text{graph}(P)$ be cycle-invariant. If \mathcal{C} guarantees authenticity, then P guarantees weak authenticity. If \mathcal{C} guarantees authenticity and P is nonce linear, then P guarantees strong authenticity.

This means that causal graphs, which express the causality among process events, are a sound (and decidable) model for proving weak authenticity while the freshness of authentication requests, namely the condition distinguishing weak from strong authenticity, may be directly verified on the syntax of ρ -spi processes. For example, since $\text{graph}(\text{Prot})$ guarantees authenticity and Prot is nonce-linear, then Prot guarantees strong authenticity. Notice that Prot describes an unbounded number of instances of A acting as claimant in protocol sessions with B and an unbounded number of instances of B acting as verifier in protocol sessions with A . We can easily extend the protocol specification with the parallel composition of A and B running protocol sessions with an arbitrary malicious party E . Even in this scenario, the protocol turns out to be safe.

6 Related Work

Causality-based modelling of concurrency is a widely studied research topic and several important results have been proposed. Event structures [32] are a general and expressive framework for modelling the causality among events in concurrent and distributed systems. This model captures the dependency among events and the interleaving of concurrent events by a partial order. A tricky problem when abstracting away from the multiplicity of protocol sessions is that an event may causally precede itself, thus loosing the antisymmetry property and, consequently, the partial order. This is easily seen by thinking of a process inputting a message and then sending out another message, which is used by the environment to construct a message sent to a replication of the former process and so on. In this scenario, since we abstract away from the multiplicity of protocol sessions, the input is causally preceded by itself. The safety of this kind of abstraction requires to determine which events abstract over the same protocol session and which ones may instead abstract over different protocol sessions. Thus a more specific structure was needed, containing some additional information about message integrity, and relying on paths representing computational flows instead of a partial order among events. Crazzola and Winskel have applied Petri nets [30], a well-known causality-based model for distributed systems, to the analysis of cryptographic protocols [14], although this work does not abstract away from the multiplicity of sessions.

Type systems proved successful in analyzing different security properties of cryptographic protocols, e.g., [1, 2]

for secrecy and [23, 12] for authenticity. As mentioned in the introduction, they exploit syntactic patterns of security protocols while providing guaranteed compositionality, but they constrain the class of analyzable protocols and are typically specific to individual security properties. On the one hand, the type system in [23] is very general and applies to several settings including authorization policies [21] and key-compromise [24], but type definitions quickly get cumbersome and no type-inference algorithm is currently available. The type system in [12] is compositional, modular, and allows for automatic type inference, but this generality is paid by restricting the form of protocols to some specific tagged patterns.

As a simple example, to the best of our knowledge, the protocol in Table 5 can be analyzed

by neither the type system from [23] nor the one from [12]. This protocol is a variation of the Needham-Schroeder-Lowe public-key authentication protocol, in which B sends a message m to A in the first message exchange. From A 's point of view, the authentication of this message is guaranteed by the handshake with nonce n_B (second and third message). The problem is that m occurs within neither the second nor the third message, thus resulting free in the corresponding type definitions and not enabling the end assertion based on nonce n_B . Furthermore, as opposed to existing type-systems that in case of failure do not yield any information on possible attacks, a simple inspection of the causal graph gives useful insights on the actual protocol run-time behaviour and, in fact, attack derivation is often immediate.

The control-flow analysis for message authenticity proposed by Bodei *et al.* in [8] and recently extended in [22] to detect replay attacks is closely related to our approach. Although the underlying static analysis techniques are different, both of the approaches rely on an abstraction of the protocol semantics and enjoy guaranteed termination. However, due to the undecidability of secrecy and authenticity, they perform an overapproximation that necessarily rules out safe protocols. Our approach relies on the causality relation among process events while the control-flow analysis of [8] statically verifies the origin and destination of messages and, more precisely, checks whether a message encrypted by A and intended for B does indeed come from A and reaches B only. A formal comparison between the tech-

niques is thus interesting but it is left as future work. Another interesting work is an abstract interpretation for mobile systems and, as an instance, for spi-calculus proposed by Feret [19]. The analysis deals with the origin of messages but does not address the freshness of authentication requests.

Strand spaces [27, 26] are an effective framework for the analysis of cryptographic protocols. Proofs of safety typically rely on the causality among protocol events and, in a recent paper [16], authors investigate how to automatically detect those specific executions, called shapes, among the infinitely many possible, that should be considered for analysis. There are interesting similarities between shapes and causal graphs and, although strand spaces do not enjoy guaranteed termination, we plan to exploit the underlying ideas to further refine the expressiveness of our analysis.


Proverif [5, 3] constitutes a powerful tool that takes as input spi-like protocol descriptions and, by Horn clause resolution, verifies a variety of different security properties such as secrecy, perfect secrecy and authenticity. The analysis is general and automated but guarantees termination only for protocols where every ciphertext is tagged differently [6].

7 Conclusion and Future Work

We have proposed a static analysis technique for analyzing security protocols based on abstract interpretation of the causality among process events. We have specifically shown that secrecy and authenticity can be soundly characterized in terms of causality, but we remark that the analysis is not tailored to these security properties but may as well be applicable to verify properties formulated in terms of causality among the actions of execution traces. The analysis enjoys guaranteed termination since the size of causal graphs is finite, the generation of paths terminates since input edges are inspected only once, thus avoiding loops due to cycles in the graph, and the analysis is linear on the number of paths. We have implemented a tool for automating the analysis, and we have applied the tool to some common protocols in the literature [18]. The analyses terminated within a few seconds and provided safety proofs for the correct versions of the protocols while failing to validate the flawed ones. Remarkably, attacks are often easily derivable by an inspection of the path sets. The only human effort required is to capture the protocol in the dialect of the spi-calculus which is often straightforwardly derivable from the protocol description.

As future work, we plan to investigate a more sophisticated abstraction allowing us to relax some of the constraints that are currently imposed by our analysis: for instance, our experiments show that some false positives occur when an authenticated term, e.g., a session key, has to be kept secret until authentication requests are accepted and is

Table 5 Variant of NSL



then leaked out. This kind of scenario turns out to be problematic also for the abstraction of Horn clauses in the decidable \mathcal{H}_1 subclass proposed by Goubault-Larrecq *et al.* [25]. A first reason for such false positives is that the leakage of names previously kept secret may introduce cycles in the graph that break the cycle-invariance condition. A possible solution is to augment the number of times each input edge can be inspected, by refining function *paths* and the cycle-invariance definition accordingly. This refinement enhances the precision of the abstraction at the price of increasing the number and size of path sets and thus the complexity of the analysis. A second reason is that the check on the freshness of nonces is used so far only as sufficient condition for proving that weak authenticity implies strong authenticity: more generally, nonce checks guarantee that different protocol sessions rely on different nonces and we believe that this information can be used for refining the analysis and, more precisely, for excluding those paths where a check on the same nonce is performed more than once, thus ruling out this kind of false positives.

For the sake of readability, we have not considered operators such as tags and hashes: their insertion in our framework does not induce any complication but is left as future work. Finally, in this paper we have only considered the instantiation of variables with names. The interesting complication arising when variables can be instantiated with ciphertexts is the capability of the environment to forge, and make trusted participants create, arbitrarily nested ciphertexts, thus potentially causing an infinite number of branches in the input. We plan to solve this problem by an overapproximation guaranteeing that the number of ciphertexts generated by trusted principals in the abstract model is finite and by abstracting away from the ciphertexts generated by the environment.

References

- [1] M. Abadi. Secrecy by typing in security protocols. *Journal of the ACM*, 46(5):749–786, 1999.
- [2] M. Abadi and B. Blanchet. Secrecy types for asymmetric communication. In *Proc. 4th International Conference on Foundations of Software Science and Computation Structures (FOSSACS)*, volume 2030 of *Lecture Notes in Computer Science*, pages 25–41. Springer-Verlag, 2001.
- [3] M. Abadi and B. Blanchet. Analyzing security protocols with secrecy types and logic programs. In *Proc. 29th Symposium on Principles of Programming Languages (POPL)*, pages 33–44. ACM Press, 2002.
- [4] M. Abadi and A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. In *Proc. 4th ACM Conference on Computer and Communications Security*, pages 36–47, 1997.
- [5] B. Blanchet. An efficient cryptographic protocol verifier based on Prolog rules. In *Proc. 14th IEEE Computer Security Foundations Workshop (CSFW)*, pages 82–96. IEEE Computer Society Press, 2001.
- [6] B. Blanchet and A. Podelski. Verification of cryptographic protocols: Tagging enforces termination. In *Proc. 6th International Conference on Foundations of Software Science and Computation Structures (FOSSACS)*, pages 136–152, 2003.
- [7] D. Bleichenbacher. Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS. In *Advances in Cryptology: CRYPTO '98*, volume 1462 of *Lecture Notes in Computer Science*, pages 1–12. Springer-Verlag, 1998.
- [8] C. Bodei, M. Buchholtz, P. Degano, F. Nielson, and H. R. Nielson. Static validation of security protocols. *Journal of Computer Security*, 13(3):347–390, 2005.
- [9] M. Boreale. Symbolic trace analysis of cryptographic protocols. In *28rd International Colloquium on Automata, Languages and Programming (ICALP 2001)*, *Lecture Notes in Computer Science*, pages 667–681. Springer-Verlag, 2001.
- [10] M. Boreale, R. D. Nicola, and R. Pugliese. Proof techniques for cryptographic processes. In *Proc. 14th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 157–166, 1999.
- [11] M. Bugliesi, R. Focardi, and M. Maffei. Compositional analysis of authentication protocols. In *Proc. 13th European Symposium on Programming (ESOP)*, volume 2986 of *Lecture Notes in Computer Science*, pages 140–154. Springer-Verlag, 2004.
- [12] M. Bugliesi, R. Focardi, and M. Maffei. Dynamic types for authentication, 2007. To appear in *Journal of Computer Security*.
- [13] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proc. 4th Symposium on Principles of Programming Languages (POPL)*, pages 238–252. ACM Press, 1977.
- [14] F. Crazzolaro and G. Winskel. Events in security protocols. In *Proc. 8th ACM Conference on Computer and Communications Security*, pages 96–105. ACM Press, 2001.
- [15] D. E. Denning and G. M. Sacco. Timestamps in key distribution protocols. *Communications of the ACM*, 24(8):533–536, 1981.
- [16] S. Doghmi, J. Guttman, and F. Thayer. Searching for shapes in cryptographic protocols. In *Proc. 13th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, *Lecture Notes in Computer Science*, pages 523–538. Springer-Verlag, 2007.
- [17] D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
- [18] L. S. et Vérification. Security protocols open repository. <http://www.lsv.ens-cachan.fr/spore/>.
- [19] J. Feret. *Analysis of Mobile Systems by Abstract Interpretation*. PhD thesis, École Polytechnique, 2005.
- [20] D. Fisher. Millions of .Net Passport accounts put at risk. *eWeek*, May 2003. (Flaw detected by Muhammad Faisal Rauf Danka).

- [21] C. Fournet, A. D. Gordon, and S. Maffei. A type discipline for authorization policies. In *Proc. 14th European Symposium on Programming (ESOP)*, Lecture Notes in Computer Science, pages 141–156. Springer-Verlag, 2005.
- [22] H. Gao, P. Degano, C. Bodei, and H. R. Nielson. Detecting replay attacks by freshness annotations. In *WITS '07: Proceedings of the 7th Workshop on Issues in the theory of security*, pages 85–100, 2007.
- [23] A. D. Gordon and A. Jeffrey. Types and effects for asymmetric cryptographic protocols. *Journal of Computer Security*, 12(3):435–484, 2004.
- [24] A. D. Gordon and A. Jeffrey. Secrecy despite compromise: Types, cryptography, and the pi-calculus. In *Proc. 16th International Conference on Concurrency Theory (CONCUR)*, volume 3653, pages 186–201. Springer-Verlag, 2005.
- [25] J. Goubault-Larrecq and F. Parrennes. Cryptographic protocol analysis on real c code. In *Proc. 6th International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI'05)*, volume 3385 of *Lecture Notes in Computer Science*, pages 363–379. Springer-Verlag, 2005.
- [26] J. Guttman, F. Thayer, J. A. Carlson, J. C. Herzog, J. D. Ramsdell, and B. T. Sniffen. Trust management in strand spaces: a rely-guarantee method. In *Proc. 13th European Symposium on Programming (ESOP)*, volume 2986 of *Lecture Notes in Computer Science*, pages 325–339. Springer-Verlag, 2004.
- [27] J. D. Guttman and F. J. Thayer. Authentication tests and the structure of bundles. *Theoretical Computer Science*, 283(2):333–380, 2002.
- [28] G. Lowe. “A Hierarchy of Authentication Specification”. In *Proc. 10th IEEE Computer Security Foundations Workshop (CSFW)*, pages 31–44. IEEE Computer Society Press, 1997.
- [29] R. Needham and M. Schroeder. Using encryption for authentication in large networks of computers. *Communications of the ACM*, 12(21):993–999, 1978.
- [30] J. L. Peterson. Petri nets. *ACM Computing Surveys*, 9(3):223–252, 1977.
- [31] D. Wagner and B. Schneier. Analysis of the SSL 3.0 protocol. In *Proc. 2nd USENIX Workshop on Electronic Commerce*, pages 29–40, 1996.
- [32] G. Winskel. Event structures. In *Advances in Petri nets 1986, part II on Petri nets: applications and relationships to other models of concurrency*, pages 325–392. Springer-Verlag, 1987.
- [33] T. Y. C. Woo and S. S. Lam. A lesson on authentication protocol design. *Operation Systems Review*, 28(3):24–37, 1994.

A Semantics of ρ -spi

The dynamics of ρ -spi is formalized by means of a transition relation, which is reported in Table 6, between *configurations*, i.e., pairs $\langle s, P \rangle$, where s is a sequence of actions and P is a closed process. Some transitions apply substitutions to processes: formally, a substitution $\sigma : x \mapsto G$ is a function from variables to run-time messages. Often substitutions are written explicitly by $[G_1/x_1, \dots, G_n/x_n]$. The

Table 6 Transition System for ρ -spi

Transition rules: We omit the symmetric rule of PAR.

NAME RES $\frac{n \notin \text{bn}(s) \cup \text{fn}(s)}{\langle s, \text{new}(n).P \rangle \rightarrow \langle s :: \text{new}(n), P \rangle}$	
KEY-PAIR RES $\frac{k_I^-, k_I^+ \notin \text{bn}(s) \cup \text{fn}(s)}{\langle s, \text{new}^\mp(k_I).P \rangle \rightarrow \langle s :: \text{new}^\mp(k_I), P \rangle}$	
INPUT $\frac{s \vdash G \quad \sigma = \text{bind}(M, G) \neq \uparrow}{\langle s, \text{in}(M).P \rangle \rightarrow \langle s :: \text{in}(M\sigma), P\sigma \rangle}$	
OUTPUT $\langle s, \text{out}(G).P \rangle \rightarrow \langle s :: \text{out}(G), P \rangle$	
BEGIN $\langle s, \text{begin}_G^i(A, I, G').P \rangle \rightarrow \langle s :: \text{begin}_G^i(A, I, G'), P \rangle$	
END $\langle s, \text{end}_G^i(A, I, G').P \rangle \rightarrow \langle s :: \text{end}_G^i(A, I, G'), P \rangle$	
PRINCIPAL $\frac{\langle s, P \rangle \rightarrow \langle s :: t, P' \rangle}{\langle s, A \triangleright P \rangle \rightarrow \langle s :: t, A \triangleright P' \rangle}$	PAR $\frac{\langle s, P \rangle \rightarrow \langle s', P' \rangle}{\langle s, P Q \rangle \rightarrow \langle s', P' Q \rangle}$
REPLICATION $\langle s, !P \rangle \rightarrow \langle s, P \mid !P \rangle$	

application of the substitution σ to the process P is denoted by $P\sigma$ and applies only to free occurrences of the variables in P . NAME RES generates a new name n by checking that it differs from all the names already used in the trace s . It is possible to force this condition by applying α -conversion to n , i.e., by substituting n and all of its free occurrences in P with a different name having the same canonical representative. Similar reasoning applies to the restriction of key-pairs. INPUT requires message G , read from the network, to be computable by the environment: the environment knowledge is defined by the message manipulation rules reported in Table 7 and discussed below. The run-time message G is read only if it can be pattern-matched with the input term M via the function bind , which is defined in Table 7 and discussed below. We write \bar{a} to denote the decryption key corresponding to a . We have $\bar{n} = n$, $\bar{k}_{IJ} = k_{IJ}$, $\bar{I} = I$, $\bar{k}_I^+ = k_I^-$ and $\bar{k}_I^- = k_I^+$. OUTPUT, BEGIN, END, and PRINCIPAL are self-explanatory. Finally, PAR interleaves two different protocol executions and REPLICATION arbitrarily replicates a principal. Function bind takes as input a static term M and a run-time message G and, in case it exists, yields the substitution σ which makes M equal to G , up

Table 7 Deduction system and binding

Message Manipulation Rules			
$\frac{\text{OUT}}{\frac{\text{out}(G) \in s}{s \vdash G}}$	$\frac{\text{ENV}}{\frac{a \notin \text{bn}(s)}{s \vdash a}}$	$\frac{\text{PAIR}}{\frac{s \vdash G_1 \quad s \vdash G_2}{s \vdash (G_1, G_2)}}$	
$\frac{\text{PAIR DES}}{\frac{s \vdash (G_1, G_2)}{s \vdash G_1 \quad s \vdash G_2}}$		$\frac{\text{ENC}}{\frac{s \vdash G \quad s \vdash a}{s \vdash \{G\}_a}}$	$\frac{\text{DEC}}{\frac{s \vdash \{G\}_a \quad s \vdash \bar{a}}{s \vdash G}}$
$\frac{\text{PUBLIC KEYS}}{s \vdash k_I^+}$		$\frac{\text{ENEMY KEYS}}{s \vdash k_{EI} \quad s \vdash k_{IE} \quad s \vdash k_E^-}$	

Binding

$$\begin{aligned}
\text{bind}(a, a) &= [] \\
\text{bind}(x, a) &= [a/x] \\
\text{bind}((M, M'), (G, G')) &= \text{bind}(M, G) \uplus \text{bind}(M', G') \\
\text{bind}(\{M\}_{\bar{a}}, \{G\}_a) &= \text{bind}(M, G) \\
\text{bind}(M, G) &= \uparrow \quad \text{otherwise}
\end{aligned}$$

to the different notation for encryption and decryption keys. If pattern-matching fails, bind returns \uparrow . This function is defined by cases on the structure of term M : a name matches a name with empty substitution; a variable can be bound to a name; pairs match pairs yielding a substitution which is the union \uplus of the ones achieved for the subterms; finally, decryptions must be performed with the correct decryption key. In all the other cases, bind returns failure \uparrow .

The knowledge of the environment is formalized by the deduction system reported in Table 7. Rule OUT says that every message sent on the network is known by the environment. ENV allows the environment to know any name which is not bound (i.e., restricted) in the trace. By PAIR and PAIR DES, the environment can construct and destruct pairs. By ENC, and DEC the environment can encrypt and decrypt messages only knowing the required keys. By PUBLIC KEYS, all the public keys are known to the environment. Finally, by ENEMY KEYS, the environment may be provided with its own private keys and with long-term keys shared with honest participants. This gives the possibility to the enemy to interact with the other participants by pretending to be a trusted principal.

B Semantics of Causal Graphs

The knowledge of the abstract environment is formalized by the judgement $\mathcal{G} \vdash G$, meaning that the environment can construct G given the knowledge of the terms in \mathcal{G} : this judgement is defined by a deduction system reported in Table 8. The environment knows every term sent on the network (OUT) and the special name \mathcal{E} (ENV), it can construct and destruct pairs (PAIR and PAIR DES) and encrypt

Table 8 Deduction System and Binding

$\frac{\text{OUT}}{\frac{G \in \mathcal{G}}{\mathcal{G} \vdash G}}$	$\frac{\text{ENV}}{\mathcal{G} \vdash \mathcal{E}}$	$\frac{\text{PAIR}}{\frac{\mathcal{G} \vdash G_1 \quad \mathcal{G} \vdash G_2}{\mathcal{G} \vdash (G_1, G_2)}}$
$\frac{\text{PAIR DES}}{\frac{\mathcal{G} \vdash (G_1, G_2)}{\mathcal{G} \vdash G_1 \quad \mathcal{G} \vdash G_2}}$		$\frac{\text{ENC}}{\frac{\mathcal{G} \vdash G \quad \mathcal{G} \vdash v}{\mathcal{G} \vdash \{G\}_v}}$
$\frac{\text{DEC}}{\frac{\mathcal{G} \vdash \{G\}_v \quad \mathcal{G} \vdash v' \quad \text{inv}(v, v')}{\mathcal{G} \vdash G}}$		

Binding

$$\begin{aligned}
\text{bind}(v_1, v_2) &= (v_2, []) && \text{if } [v_1] = [v_2] \vee \\
&&& ([v_1] = \mathcal{E} \wedge [v_2] \in PN) \vee \\
&&& ([v_2] = \mathcal{E} \wedge [v_1] \in PN) \\
\text{bind}(x, v) &= (v_{(x)}, [v_{(x)}/x]) \\
\text{bind}((M, M'), (G, G')) &= && \text{if } \text{bind}(M, G) = (G_1, \sigma_1) \wedge \\
&&& ((G_1, G_2), (\sigma_1 \uplus \sigma_2)) && \text{bind}(M', G') = (G_2, \sigma_2) \\
\text{bind}(\{M\}_v, \{G\}_{v'}) &= (\{G'\}_{v'}, \sigma) && \text{if } \text{bind}(M, G) = (G', \sigma) \wedge \\
&&& \text{inv}(v, v') \\
\text{bind}(M, G) &= \uparrow && \text{otherwise}
\end{aligned}$$

Notation:

$$\begin{aligned}
PN &\triangleq \{a \mid \exists k, l, E \text{ s.t. } a \in \{l, k_I^+, k_E^-, k_{IE}, k_{EI}, \mathcal{E}\}\}. \\
\text{In OUT, PAIR DES, and DEC, } G, G_1, \text{ and } G_2 &\text{ are not in } PN. \\
\text{inv}(v, v') &\triangleq ([v'] \neq \mathcal{E} \Rightarrow \text{bind}(v, [v']) \neq \uparrow) \wedge \\
&([v] \neq \mathcal{E} \Rightarrow \text{bind}([v], v') \neq \uparrow)
\end{aligned}$$

and decrypt terms only knowing the required keys (ENC and DEC). For reducing the number of terms known to the environment and abstracting the same set of ρ -spi terms, we prevent the environment from deriving labelled version of identities, public keys, enemy's keys and \mathcal{E} , which abstract the same ρ -spi terms as \mathcal{E} .

Function bind , reported in Table 8, defines the pattern-matching among terms. This function takes as input a term M and a ground term G and, if the two terms match, yields the term G' obtained by labelling G according to the variables in M and the corresponding substitution. If pattern-matching fails, bind returns \uparrow . Function bind is defined by cases on the structure of term M : a name matches itself and, similarly, the name \mathcal{E} matches identities, public keys and enemy's keys; a variable can only be bound to an atomic name v ; pattern-matching of pairs and ciphertexts is defined component-wise: notice that decryptions must be performed by the correct decryption key. In all the other cases, bind returns failure \uparrow . Function \bar{v} yields the decryption key matching the encryption key v : this function is smoothly extended to arbitrary abstract terms G by inverting the encryption keys in G .