

Causality-based Abstraction of Multiplicity in Security Protocols

Michael Backes¹, Matteo Maffei¹, and Agostino Cortesi²

¹ Saarland University, Saarbrücken, Germany, {backes,maffei}@cs.uni-sb.de

² 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, 30], 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 for these patterns. 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 [22, 11]. This rules out some interesting protocols relying on composing multiple authentication sessions based on the same nonce, e.g., see [25].

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 a *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 agreement and injective agreement [26].

We can summarize our technique as follows. We first specify the protocol in a dialect of the spi-calculus [3]. 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 techniques 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 some common protocols in the literature [18] (among others, the Needham-Schroeder-Lowe public-key protocol, its fixed version proposed by Lowe, the BAN modified version of CCITT X.509(3) and SPLICE/AS). 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 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.

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

Table 1 (Our Dialect of) the ρ -spi Calculus**Notation:** u ranges over names and variables.

Name	Process
$a ::= I, J, A, B, E$ (Id)	$P, Q ::= \text{new}(n).P$ (Msg Res)
n, m, k_{IJ} (Msg)	$\text{new}^\mp(k_I).P$ (Key Res)
k_I^+ (Pub)	$\text{in}(M).P$ (In)
k_I^- (Priv)	$\text{out}(M).P$ (Out)
Term	$\text{begin}_N^i(A, I, M).P$ (Begin)
$M, N ::= a$ (Name)	$\text{end}_N^i(A, I, M).P$ (End)
x, y, z (Var)	$A \triangleright P$ (Principal)
(M, N) (Pair)	$P Q$ (Par)
$\{M\}_u$ (Enc)	$!P$ (Repl)
	$\mathbf{0}$ (Stop)

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] is derived from the spi calculus [3] 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 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 of our dialect of the ρ -spi calculus

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 *messages* \mathcal{M} , the set of identities \mathcal{ID} , the set of shared keys \mathcal{K} and the set of public and private keys \mathcal{K}^\mp . The set \mathcal{ID} , ranged over by I and J , is further partitioned into the two sets of *trusted principals* \mathcal{ID}_P , ranged over by A and B , and *enemies* \mathcal{ID}_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}_E contains the keys of malicious parties, i.e., the ones where some 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.³

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 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, 24], we assume that messages are typed so as to distinguish names from the other terms like pairs and ciphertexts. In the long version of this paper, we extend our approach to additionally allow variables to be instantiated with ciphertexts.⁴ The interesting complication arising with such an extension is the capability of the environment to forge arbitrarily nested ciphertexts, thus potentially causing an infinite number of branches in the input. We tackle this problem in the abstract model by guaranteeing that the number of ciphertexts generated by trusted principals is finite and by abstracting away from the ones generated by the environment.


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. To illustrate, let us consider the following protocol, in which B encrypts the message m , the nonce n and his own identifier with A ’s public-key and A acknowledges

³ For the sake of readability, in the rest of the paper we omit brackets: for instance, the nested pair $((a, b), k)$ is simplified in a, b, k .

⁴ For convenience, this is also shown in Appendix D.

the reception of the first message by sending back the nonce in clear on the network.



The begin assertion says that A confirms the reception from B of the 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 strong authenticity between A and B .

2.2 Operational semantics of our dialect of the ρ -spi calculus

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 *agreement* property of [26, 33] by making explicit the nonce used in the handshake and the index binding the begin assertion with the corresponding end one.

Definition 3 (Weak Authenticity). *A trace s guarantees weak authenticity if and only if $s = s_1 :: \text{end}_{G_1}^1(B, A, G_2) :: s_2$ implies $\text{begin}_{G_1}^1(A, B, G_2) \in s_1$. 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 *injective agreement* property of [26, 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 if 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 the 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*, and we demand that two messages are α -convertible only when they have the same canonical representative. A similar assumption applies to key-pairs and variables. We write $\llbracket n \rrbracket$ to denote the set of messages with the same canonical representative as n . The notation is extended to key-pairs and variables as expected. We also require that keys, both symmetric and asymmetric ones, having the same canonical representative depend on the same identifiers. We assume that the set \mathcal{M}_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. Finally, we shall write $\llbracket P \rrbracket$ to denote the set of processes obtained from P by α -conversion of bound names and variables.

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 are divided into *process nodes*, representing abstraction of ρ -spi processes, and *communication nodes*, representing the synchronization among a process performing an input and the processes outputting the terms used by the environment to produce the input term. Edges track the causality among nodes. In the following, we discuss 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:

- *the special name* \mathcal{E} : it arises in the terms forged by the environment and abstracts the names generated by the environment, identities, public-keys and attackers' keys;
- *labelled names* $a_{(x)}$: they are semantically equivalent to names but they 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 in 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 \mathcal{G} over ground term sets.


Edges connect two nodes, 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 $P \rightarrow Q$, directly linking P to the process Q obtained by reducing P .

Inter-thread causality is due to the synchronization among a process P inputting G and reducing to Q and the nodes $\text{out}(G_i).P_i$ outputting the terms used by the environment to construct the input term 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 the 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 $\text{in}(M).P \text{---} \boxed{\text{com}(G)} \rightarrow Q$, connecting the input process with the process obtained by performing the input via the synchronization point $\text{com}(G)$; and (ii) output edges of the form $\text{out}(G_i) \xrightarrow{G_i} \boxed{\text{com}(G)} \rightarrow Q$, 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 $\text{com}(G)$.

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:



Notice that the terms in the communication node track both the positional indexes of the names occurring therein 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 $[M]$ to denote the term obtained from M by label erasure and $\lceil M \rceil$ to denote the term obtained from M by index erasure.

Function $output(out(G).P)$ yields the term obtained from G by (i) erasing the labels in G and (ii) indexing each name a in G with the pair composed of process $out(G).P$ 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, $output(\mathcal{N}) = \{output(out(G).P) \mid out(G).P \in \mathcal{N}\}$. We write $index(G)$ to denote the set of processes occurring in the indexes of G . 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} . The function $bind_{abs} : (M, G) \mapsto G'$, 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 . The substitution σ expressing the variable instantiation induced by the pattern-matching of M with G is returned by the function $subst(G')$, with $G' = bind_{abs}(M, G)$: the function $subst : G \mapsto \sigma$ yields the least substitution σ such that for every $v_{(x)} \in names(G)$, $\sigma(x) = \lceil v \rceil_{(x)}$. If pattern-matching fails, $bind_{abs}$ and consequently $subst$ return \uparrow . For example, $bind_{abs}(x, n^5) = n^5_{(x)}$ and $subst(n^5_{(x)}) = [n_{(x)}/x]$. Function $threads$, defined below, yields the set of threads in P , abstracting away from identifiers and replications.

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

Finally, function $int(\mathcal{N}, G)$ defines the set of integer terms in G : these are defined as the largest subterms of G that have not been generated by the environment, i.e., either names or ciphertexts whose encryption key is unknown to the environment.

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

The next definition characterizes the causal graph associated with an abstract process. Here and throughout this paper, we let $\llbracket P \rrbracket$ denote the set of processes obtained from P by α -renaming of bound names and $\llbracket P \rrbracket \approx$ denote the set of processes obtained from P by α -renaming of bound names *and* by applying substitutions of the form $[n_{(x)}/m_{(x)}]$, for some $m \in \llbracket n \rrbracket$.

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

$$\begin{aligned} & threads(P) \subseteq \mathcal{N} && \text{(Initial)} \\ & p.P \in \mathcal{N} \wedge Q \in threads(P) \wedge p \neq in(\cdot) && \text{(Intra-thread)} \\ & \Rightarrow Q \in \mathcal{N} \wedge p.P \rightarrow Q \in \mathcal{E} \\ & in(M).P \in \mathcal{N} \wedge output(\mathcal{N}) \vdash G \wedge bind_{abs}(M, G) = G' \wedge \sigma = subst(G') \wedge && \text{(Inter-thread)} \\ & \quad P' \in threads(P\sigma) \wedge G'_1 \in int(\mathcal{N}, G') \wedge \{out(G_1).P_1\} = index(G'_1) \\ & \Rightarrow \exists Q \in \llbracket P' \rrbracket \text{ s.t. } Q \in \mathcal{N} \wedge \{in(M).P \xrightarrow{\text{com}(G')} \rightarrow Q, out(G_1).P_1 \xrightarrow{G'_1} \text{com}(G') \rightarrow Q\} \subseteq \mathcal{E} \\ & \forall \{new(n).P, new(m).Q\} \subseteq \mathcal{N}, \quad new(n).P \in \llbracket new(m).Q \rrbracket \approx \Leftrightarrow n = m && \text{(Uniqueness)} \end{aligned}$$

Function $graph : P \mapsto (\mathcal{N}, \mathcal{E})$ is in fact a closure operator on causal graphs ordered by inclusion. In particular, the threads in P are part of the set \mathcal{N} of nodes (*Initial*); if $p.P$ is a node in the causal graph, with $p \neq in(\cdot)$, then the processes in $threads(P)$, which are connected to $p.P$ by direct edges, belong to the set of nodes (*Intra-thread*); if $in(M).P$ is a node and the environment knows a term G matching M with substitution σ , then the set of nodes contains also the processes in $threads(P\sigma)$ and the communication node linking the output processes used by the environment to construct G to the processes in $threads(P\sigma)$ (*Inter-thread*). 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. Finally, notice that condition (*Inter-thread*) allows for α -renaming the bound names in the processes following an input ($\exists Q \in \llbracket P' \rrbracket \dots$). This α -renaming allows the abstract domain to distinguish the names generated in different threads, thus making the analysis more precise. Unfortunately, this might introduce an infinite number of nodes in the causal graph, which only differ because of α -renaming. The solution is to α -rename bound names whenever the input term is different from the ones already considered, where the difference among input terms is up to α -renaming of the names bound to variables: condition (*Uniqueness*) states that the causal graph does not contain copies of the same node differing only because of α -renaming of bound names, thus 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 $length(P)$ to denote the number of primitives in P and $length(\mathcal{N})$ to denote the maximal length of the processes in \mathcal{N} , namely the number n such that $length(Q) = n$, for some $Q \in \mathcal{N}$, and, for every $Q' \in \mathcal{N}$, $length(Q') \leq n$.

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 $length(\mathcal{N}) = length(\mathcal{N}')$. 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 *Inter-thread*). However, because of condition (*Uniqueness*), the number of α -renamed copies of name n 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

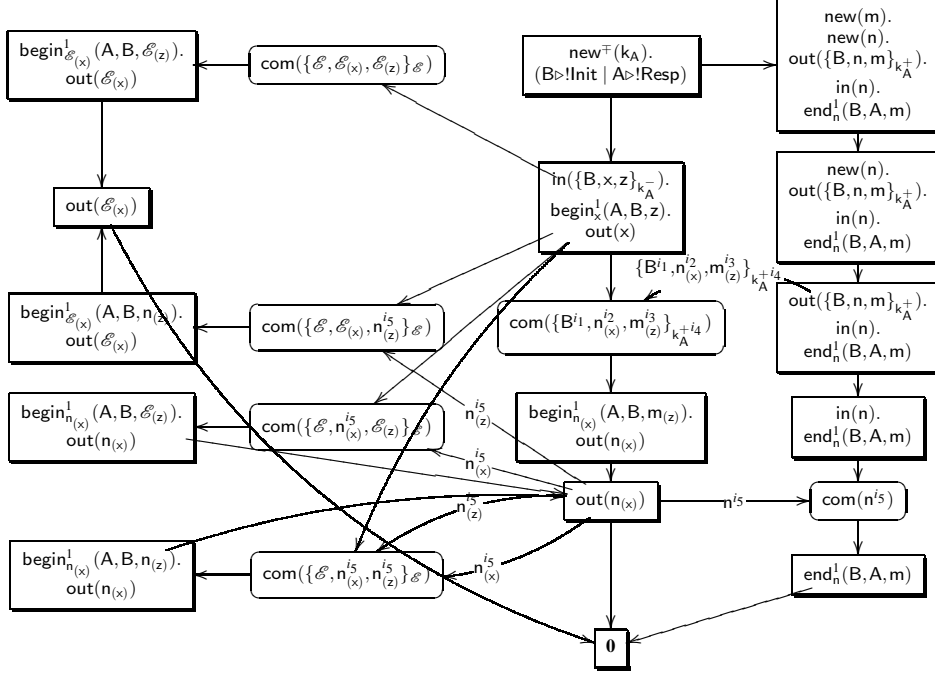
The following corollary says that the size of the graph grows exponentially with the protocol *specification*, which is however fixed in advance, regardless of the number of considered sessions and the protocol run-time behavior. We remark that this result refers to the worst case: in practice, the number of nodes for the protocols we have considered so far does not exceed 100 nodes and the analysis always terminates in less than 1 second.¹

¹. This has to be verified!

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

Example 3. To illustrate, the causal graph associated with protocol *Prot* of Example 1 is depicted in Table 2. The rounded boxes represent communication nodes while the other ones denote process nodes. 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 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

Table 2 Causal graph $graph(Prot)$



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

in the causal graph: we can see that $output(\mathcal{N}) \not\vdash m^i$ 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 $begin_{n(x)}^1(A, B, \mathcal{E}(z)).out(n_{(x)})$ and $begin_{n(x)}^1(A, B, n_{(z)}).out(n_{(x)})$. These events are enabled by the environment forging the messages $\{\mathcal{E}, n^{i5}, \mathcal{E}\}_{\mathcal{E}}$ and $\{\mathcal{E}, n^{i5}, n^{i5}\}_{\mathcal{E}}$, respectively, 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 messages in which the message \mathcal{E} replaces the nonce, but the resulting processes do not increase the environment’s knowledge as they eventually output \mathcal{E} , which is always 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 $\text{out}(n).P \rightarrow P$, then the process P is causally preceded by the reduction of the process $\text{out}(n).P$ with action $\text{out}(n)$. In general, we express the causality dependency among process events through *paths*, which are sequences of *actions* expressing process events:

$$\begin{aligned} t &::= \text{new}(n) \mid \text{new}^\mp(k_1) \mid \text{in}(G) \mid \text{out}(G) \mid \text{in}_{\text{com}}(G, G) \mid \text{begin}_G^i(A, I, G) \mid \text{end}_G^i(A, I, G) \\ s &::= \varepsilon \mid s :: t \end{aligned}$$

Abstract actions extend ρ -spi ones with the new action $\text{in}_{\text{com}}(G_1, G_2)$, which is used for tracking inputs involved in inter-thread causality: G_1 is the integer component and G_2 is the input term. 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 construct 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 edges connecting an input node with a communication one only once, thus abstracting away from cycles in the causal graph which, in fact, do not alter the causality among nodes.

We write $\{s_1, \dots, s_n\} :: \tau$ to denote $\{s_1 :: \tau, \dots, s_n :: \tau\}$ and $\text{outedges}(G, Q, \mathcal{C})$ to denote the set of output edges in \mathcal{C} incoming in Q via the synchronization point $\text{com}(G)$. The function $\text{paths} : \mathcal{C}, P, \mathcal{E} \mapsto \mathcal{S}$, 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 $\text{paths}(\mathcal{C}, P)$ to denote $\text{paths}(\mathcal{C}, P, \emptyset)$.

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

$$\begin{aligned} Q \text{ has no incoming edge} &\Rightarrow \mathcal{S} = \{\emptyset\} && \text{(Initial)} \\ p.P \rightarrow Q \in \text{edges}(\mathcal{C}) \wedge \mathcal{S} \in \text{paths}(\mathcal{C}, p.P, \mathcal{E}) &\Rightarrow \mathcal{S} :: p \in \mathcal{S} && \text{(Intra)} \\ \text{in}(M).P \text{---} \boxed{\text{com}(G)} \rightarrow Q \in \text{edges}(\mathcal{C}) \setminus \mathcal{E} \wedge \sigma = \text{subst}(G) \wedge && \text{(Inter)} \\ \{\text{out}(G_1).P_1 \xrightarrow{G_1} \boxed{\text{com}(G)} \rightarrow Q, \dots, \text{out}(G_n).P_n \xrightarrow{G_n} \boxed{\text{com}(G)} \rightarrow Q\} = \text{outedges}(G, Q, \mathcal{C}) \wedge && \\ \mathcal{S} \in \text{paths}(\mathcal{C}, \text{in}(M).P, \mathcal{E} \cup \{\text{in}(M).P \text{---} \boxed{\text{com}(G)} \rightarrow Q\}) \wedge && \\ \mathcal{S}_i \in \text{paths}(\mathcal{C}, \text{out}(G_i).P_i, \mathcal{E} \cup \{\text{in}(M).P \text{---} \boxed{\text{com}(G)} \rightarrow Q\}) && \\ \Rightarrow \mathcal{S} :: \text{in}(M\sigma) \cup \bigcup_{i \in [1, n]} \mathcal{S}_i :: \text{out}(G_i) :: \text{in}_{\text{com}}(G'_i, M\sigma) \in \mathcal{S} && \end{aligned}$$

If the node Q is preceded by intra-thread causality by $p.P$, 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 a process receiving G ($\text{in}(M).P \xrightarrow{\text{com}(G)} Q$), where the input edge $\text{in}(M).P \xrightarrow{\text{com}(G)} Q$ is not in \mathcal{E} , then the paths associated with Q are the ones associated to $\text{in}(M).P$ ($\mathcal{S} \in \text{paths}(\mathcal{C}, \text{in}(M).P, \mathcal{E} \cup \{\text{in}(M).P \xrightarrow{\text{com}(G)} Q\})$), each of them extended with the action $\text{in}(M\sigma)$, plus the ones associated with the processes outputting the terms used to construct term G ($\{\text{out}(G_1).P_1 \xrightarrow{G'_1} \text{com}(G) \rightarrow Q, \dots, \text{out}(G_n).P_n \xrightarrow{G'_n} \text{com}(G) \rightarrow Q\} = \text{outedges}(G, \mathcal{C})$ and $\mathcal{S}_i \in \text{paths}(\mathcal{C}, \text{out}(G_i).P_i, \mathcal{E} \cup \{\text{in}(M).P \xrightarrow{\text{com}(G)} Q\})$), extended with $\text{out}(G'_i) :: \text{in}(G_i, M\sigma)$. Notice that the edge $\text{in}(M).P \xrightarrow{\text{com}(G)} Q$ is inspected only once, thus avoiding loops due to cycles in the graph.

Example 4. For instance, $\text{paths}(\text{graph}(\text{Prot}), \text{end}_n^1(B, A, m), \emptyset)$, where $\text{graph}(\text{Prot})$ is the causal graph in Table 2, yields the path sets reported in Table 9 of Appendix B. The number of possible paths depends on the environment's capability of forging challenges and thus interacting with the responder. Notice that every path set contains a path whose prefix is as follows:

$$c \triangleq \text{new}^\mp(k_A) :: \text{new}(n) :: \text{new}(m) :: \text{out}(\{B, n, m\}_{k_A^+}) :: \text{in}_{\text{com}}(\{B^{i_1}, n_{(x)}^{i_2}, m_{(z)}^{i_3}\}_{k_A^{+i_4}}, \{B, n_{(x)}, m_{(z)}\}_{k_A^-}) :: \text{begin}_{n_{(x)}}^1(A, B, m_{(z)})$$

Notice also that this path represents the intended protocol behavior. Intuitively, this means that the environment can play a number of actions but these necessarily follow the intended protocol run and thus they do not affect the safety of the protocol as far as authenticity is concerned.

3.4 Names and sessions

Since a causal graph abstracts an unbounded number of protocol sessions, an interesting issue related to our abstraction is that different occurrences in a path of the same name n might actually abstract different ρ -spi names, one for each 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 protocol security goals. According to the definition of causal graphs, the labels of the edges incoming in communication nodes express the integer components in 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. Indeed, 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 have been generated in different protocol sessions. Here and throughout this paper, we write $\text{msgs}(M)$ 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 . Remember that a path

Table 3 Equality among abstract terms

$\frac{\text{INTRA} \quad \text{prev}(s, v) = j \quad s = i}{s \vdash (v, i) = (v, j)}$	$\frac{\text{INTER-BIND} \quad v^{(P, j)} \in \text{msgs}(G_2) \quad \pi^j(G_1) = v' \quad s = i}{s :: \text{out}(G_1) :: \text{incom}(G_2, G) \vdash (v, i+2) = (v', i+1)}$	
$\text{INTER-MATCH} \quad \frac{s :: t = s_1 :: \text{out}(G_1) :: \text{incom}(G_2, G) :: s_2 \quad \text{incom}(\dots, \dots) \notin s_2 \quad n \in \text{msgs}(G) \quad s_1 = s'_1 :: \text{new}(n) :: s'_2 \quad s = i \quad s_1 = j \quad s'_1 = j' \quad s \vdash (n, j+2) = (n, j'+1) \quad \text{prev}(s'_1, v) = k}{s \vdash (v, i) = (v, k)}$		
$\frac{\text{TRANS} \quad s \vdash (v, i) = (v', i') \quad s \vdash (v', i') = (v'', i'')}{s \vdash (v, i) = (v'', i'')}$	$\frac{\text{SYM} \quad s \vdash (v, i) = (v', i')}{s \vdash (v', i') = (v, i)}$	$\frac{\text{PREFIX} \quad s \vdash (v, i) = (v', i')}{s :: t \vdash (v, i) = (v', i')}$

may “cross” different threads: in the following, we say that an action in a path s belongs to the last thread of s if it occurs within a suffix of s that does not contain any action of the form $\text{incom}(G_1, G_2)$. Finally, function $\text{prev} : (s, v) \mapsto i$ takes as input a path s and a possibly labelled message v and yields either the position i of the last action in the last thread of s containing message v or \uparrow , if such an action does not exist.

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 consists in 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 INTRA says that the occurrences of a message within actions related by intra-thread causality abstract over the same ρ -spi name: actions related by intra-thread causality clearly 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 the 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 and SYM make the relation transitive and symmetric and PREFIX makes it closed by prefixes.

Example 5. Let us consider the following path, obtained by completing the path c of Example 4 so as to get the intended protocol behavior:

$$\begin{aligned} d \triangleq & \text{new}^\mp(k_A) :: \text{new}(m) :: \text{new}(n) :: \text{out}(\{B, n, m\}_{k_A^+}) :: \\ & \text{incom}(\{B^{i_1}, n_{(x)}^{i_2}, m_{(z)}^{i_3}\}_{k_A^{+i_4}}, \{B, n_{(x)}, m_{(z)}\}_{k_A^-}) :: \text{begin}_{n_{(x)}}^1(A, B, m_{(z)}) \\ & \text{out}(n_{(x)}) :: \text{incom}(n^{i_5}, n) :: \text{end}_n^i(B, A, m) \end{aligned}$$

The messages occurring in the begin and end assertions abstract over the same protocol session. In particular, we can prove $d \vdash (n_{(x)}, 6) = (n, 9)$ via INTER-BIND and

$d \vdash (m_{(z)}, 6) = (m, 9)$ via INTER-MATCH ($d \vdash (m, 2) = (m, 9)$), INTER-BIND ($d \vdash (m_{(z)}, 6) = (m, 2)$) and TRANS.

In Section 3.3, we have defined the paths associated with a node in a causal graph and, for avoiding loops due to cycles in the graph, we have requested that edges connecting an input node with a communication one 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 any path.

Definition 7 (Cycle-invariance). *We say that \mathcal{C} is cycle-invariant if and only if whenever $\{\text{in}(M).P \text{---} \boxed{\text{com}(G)} \rightarrow Q, \text{out}(G_1).P' \xrightarrow{G_2} \boxed{\text{com}(G)} \rightarrow Q\} \subseteq \text{edges}(\mathcal{C})$, the following conditions hold:*

- for every $S :: \text{in}(G') \in \text{paths}(\mathcal{C}, Q, \emptyset)$, $S' \in \text{paths}(\mathcal{C}, \text{in}(M).P, \emptyset)$ and $s_1 :: \text{in}(G') :: s_2 \in S'$, there exists $s = s_1 :: \text{in}(G') :: s_2' \in S'$ such that $s_1 \vdash (v, j) = (v', j')$ implies $(s) \vdash (v, i) = (v', j')$, with $|s| = i$ and $|s_1| = j$.
- for every $S :: \text{out}(G_1) :: \text{in}_{\text{com}}(G_2, G') \in \text{paths}(\mathcal{C}, Q, \emptyset)$, $S' \in \text{paths}(\mathcal{C}, \text{out}(G_1).P', \emptyset)$ and $s_1 :: \text{out}(G_1) :: \text{in}_{\text{com}}(G_2, G') :: s_2 \in S'$, there exists $s = s_1 :: \text{out}(G_1) :: \text{in}_{\text{com}}(G_2, G') :: s_2' \in S'$ such that $s_1 \vdash (v, j) = (v', j')$ and $v \in \text{msgs}(G_1)$ implies $s \vdash (v, i) = (v', j')$, with $|s| = i$ and $|s_1| = j$.

For example, $\text{graph}(\text{Prot})$ 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 preserve $n_{(x)}$ do not affect the equality constraints. In our experiments, we did not find any protocol whose specification is not cycle-invariant.²

². To be checked

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 $\gamma_{\text{term}} : M \mapsto \{M_1, \dots, M_n\}$. Abstract identities are instantiated by 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

Table 4 Concretization of paths, processes and causal graphs

$$\gamma_{\text{trm}}(M) \triangleq \begin{cases} \{I\} & \text{if } M = I \\ \{n \mid n \in \llbracket m \rrbracket\} & \text{if } M = m \\ \{y \mid y \in \llbracket x \rrbracket\} & \text{if } M = x \\ \{k_I^+ \mid k_I^+ \in \llbracket k_I \rrbracket\} & \text{if } M = k_I^+ \\ \{k_I^- \mid k_I^- \in \llbracket k_I \rrbracket\} & \text{if } M = k_I^- \\ \mathcal{ID} \cup \mathcal{K}^+ \cup \mathcal{M}_E \cup \mathcal{K}_E & \text{if } M = \mathcal{E} \\ \gamma_{\text{trm}}(a) & \text{if } M \in \{a_{(x)}, a^i, a^i_{(x)}\} \\ \{M\sigma \mid \forall u \in \text{dom}(\sigma), \sigma(u) \in \gamma_{\text{trm}}(u)\} & \text{otherwise} \end{cases}$$

$$\begin{aligned} \gamma_{\text{path}}(s) &\triangleq \{(s, \sigma) \mid |s| = |\sigma| = k \wedge \\ &\quad (i) \forall i \in [1, k], \exists \sigma_i \text{ s.t. } \pi^i(\bar{s}) = \pi^i(s)\sigma_i \wedge s \vdash (v, j) = (v', j') \Rightarrow \sigma_j(v) = \sigma_{j'}(v') \\ &\quad (ii) s \vdash (v, k) = (v', k') \wedge \sigma_k(v') = a \Leftrightarrow \sigma(v) = a \} \end{aligned}$$

$$\gamma_{\text{pset}}(\mathcal{S}) \triangleq \{(\{s_1, \dots, s_n\}, \sigma) \mid \mathcal{S} = \{s_1, \dots, s_n\} \wedge \forall i \in [1, n], (s_i, \sigma_i) \in \gamma_{\text{path}}(s_i) \wedge \sigma = \bigsqcup_{i=1, n} \sigma_i \neq \uparrow\}$$

$$\gamma_{\text{net}}(\mathcal{C}) \triangleq \{(t_1 :: \dots :: t_n, P) \mid$$

$$\begin{aligned} &\quad (i) \forall i \in [1, n], \exists Q_i \in \text{nodes}(\mathcal{C}), S_i \in \text{paths}(\mathcal{C}, Q_i, \emptyset), (S_i, \sigma_i) \in \gamma_{\text{pset}}(S_i) \text{ s.t. } S_i = S_j :: t_i, \text{ with } j < i \\ &\quad (ii) \forall Q \in \text{threads}(P), \exists i \in [1, n], \sigma \text{ s.t. } Q = Q_i \sigma \wedge \sigma \sqcup \sigma_i \neq \uparrow \} \end{aligned}$$

Notation:

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

\bar{s} yields the trace obtained from s by replacing each occurrence of $\text{incom}(\mathbf{G}_1, \mathbf{G}_2)$ by $\text{in}(\mathbf{G}_2)$

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

$$\sigma_1 \sqcup \sigma_2 = \uparrow \quad \text{otherwise}$$

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. We recall that a path may “cross” different threads and $s \vdash (v, k) = (v', k')$, with k being the length of s , only if v is an abstract name occurring the last thread.

3. To be said somewhere

Example 6 (Path Concretization). To illustrate, a concretization of the path d reported in Example 5 is as follows:

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

Since $d \vdash (n_{(x)}, 6) = (n, 9)$, $n_{(x)}$ and n are instantiated with the same name n . Similarly, since $d \vdash (m_{(z)}, 6) = (m, 9)$, 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 ($\biguplus_{i=1,n} \sigma_i \neq \uparrow$): the union $\sigma_1 \uplus \sigma_2$ 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 \mathcal{S}_i associated with Q_i and a trace set S_i such that S_i is an instantiation of \mathcal{S}_i ($(S_i, \sigma_i) \in \gamma_{\text{pset}}(\mathcal{S}_i)$), where S_i is a trace set associated with 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 \mathcal{S}_i , a trace set S_i and an instantiation σ_i such that $Q = Q_i \sigma$, where σ is compatible with the instantiation σ_i of the names 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 with 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 $\langle \varepsilon, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$.*

Due to space constraints, we postpone the proofs to Appendix 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). *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 $\text{output}(\text{nodes}(\mathcal{C})) \vdash v'$ implies $\lceil v \rceil \neq \lceil v' \rceil$.*

As stated below, a path guarantees authenticity if (i) every end assertion is preceded by a corresponding begin assertion; and (ii) begin and end assertions contain only session messages. 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 $\mathcal{S} \in \text{paths}(\mathcal{C}, \text{end}_{G_1}^i(A, B, G_2).P)$, there exists a path in \mathcal{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 $\lfloor G \rfloor = \lfloor G' \rfloor$ (i.e., the two terms are equal up to labelling) and for every k such that $\pi^k(G) = v$, $\pi^k(G') = v'$ and v, v' are possibly labelled messages, we have that $s \vdash (v, i) = (v', j)$ (i.e., the two terms abstract the same ρ -spi term).

Definition 9 (Abstract Authenticity). *A path s guarantees authenticity if and only if for every i such that $\pi^i(s) = \text{end}_{G_1}^i(A, B, G_2)$, the following conditions hold:*

i there exist $G'_1, G'_2, i < j$ s.t. $\pi^j(s_1) = \text{begin}_{G'_1}^i(B, A, G'_2)$, $G_1 \simeq_{s,i,j} G'_1$ and $G_2 \simeq_{s,i,j} G'_2$

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

For example, the path d in Example 5 guarantees authenticity as the end assertion is preceded by a corresponding begin assertion and the messages occurring therein are equivalent. By an inspection of Table 9, every $\mathcal{S} \in \text{paths}(\text{graph}(\text{Prot}), \text{end}_n^1(B, A, m))$ 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 $v \in \gamma_{\text{tm}}(v)$.*

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 $\text{graph}(Prot)$ 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 losing 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. Crazzolaro and Winskel have applied Petri nets [29], a well-known causality-based model for distributed systems, to the analysis of cryptographic protocols [14] but their work does not abstract away from the multiplicity of sessions. However, Petri nets might constitute a suitable alternative foundation for our work.

Type systems proved successful in analyzing different security properties of cryptographic protocols, e.g., [1, 2] for secrecy and [22, 11] 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 [22] is very general and applies to several settings including authorization policies [21] and key-compromise [23], but type definitions quickly get cumbersome and no type-inference algorithm is currently available. The type system in [11] 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 [22] nor the one from [11]. 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 . Another remarkable difference with respect to

existing type-based approaches, which in case of failure do not give any information on possible attacks, is that the analysis of computational flows does yield insights on possible attacks and, in fact, their derivation is often immediate.

The work by Bodei *et al.* on control-flow analysis for message authentication in *Lysa* [8] is closely related to our approach. The motivations and goals are different though, since the analysis only applies to message authenticity, namely the origin of a message, without addressing freshness of authentication requests and the willingness of the claimant to authenticate with the verifier. Another interesting work is an abstract interpretation for mobile systems and, as an instance, for spi-calculus [19]. The analysis deals with the origin of messages but does not address freshness.


Strand spaces [24, 25] constitute an effective framework for studying secrecy and authenticity. While some effective patterns for deriving the proofs of safety are provided, these are in general written by hand. Interestingly, such proofs often rely on the causality among events and it would be interesting to explore how our analysis could be applied in such a framework to provide mechanical proofs of safety. A recent paper [16] investigates how to automatically detect those specific executions, among the infinitely many possible, that should be considered for analysis, by identifying causality between nodes in graphs and describing execution of protocol roles by different processes. Although this method does not enjoy guaranteed termination, we plan to exploit the ideas underlying this result to further refine the expressivity of our analysis.

Proverif [4, 5] 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 fully automated but guarantees termination only for protocols where every ciphertext is tagged differently [6].

7 Conclusion and Future Work

We have proposed a novel technique for analyzing security protocols based on abstract interpretation of the causality among process events. In this paper, 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; we are currently extending our technique to analyze more sophisticated security properties emerging in modern applications such as e-voting protocols. The analysis enjoys guaranteed termination since the size of causal graphs is finite, the generation of paths terminates since communication edges (i.e., edges connecting an input node with a communication one) 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

Table 5 Variant of NSL



common protocols in the literature [18] (among others, the Needham-Schroeder-Lowe public-key protocol, its fixed version proposed by Lowe, the BAN modified version of CCITT X.509(3) and SPLICE/AS). 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 then leaked out. This may be exploited for modeling session key corruption. The reason for such false positives is that the check on the freshness of nonces is used so far 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 in causal graph where a check on the same nonce is performed more than once, thus ruling out this kind of false positives.

We remark that, 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, we plan to investigate compositionality properties on causal graphs and to apply the principles underlying our technique to the applied pi-calculus, whose flexible and general equational theory for terms will give rise to interesting issues concerning the convergence of the analysis.

References

1. M. Abadi. Secrecy by typing in security protocols. *JACM*, 46(5):749–786, 1999.
2. M. Abadi and B. Blanchet. Secrecy types for asymmetric communication. *Theoretical Computer Science*, 298(3):387–415, 2003.
3. M. Abadi and A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148(1):1–70, 1999.
4. B. Blanchet. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *Proceedings of 14th IEEE Computer Security Foundations Workshop (CSFW-14)*, pages 82–96, Cape Breton, Nova Scotia, Canada, June 2001. IEEE Computer Society.
5. B. Blanchet. From Secrecy to Authenticity in Security Protocols. In Manuel Hermenegildo and Germán Puebla, editors, *9th International Static Analysis Symposium (SAS'02)*, volume 2477 of *Lecture Notes on Computer Science*, pages 342–359, Madrid, Spain, September 2002. Springer Verlag.
6. B. Blanchet and A. Podelski. Verification of cryptographic protocols: Tagging enforces termination. In *Proceedings of Foundations of Software Science and Computation Structures*, 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 *LNCS*, 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 *Proceedings of 28th International Colloquium on Automata, Languages and Programming (ICALP 01)*, volume 2076, pages 667–681. LNCS 2076, Springer Verlag, 2001.
10. M. Boreale, R. De Nicola, and R. Pugliese. Proof techniques for cryptographic processes. In *Proceedings of 14th Symposium on Logic in Computer Science (LICS '99)*, pages 157–166, 1999.
11. M. Bugliesi, R. Focardi, and M. Maffei. Dynamic types for authentication. To appear in *Journal of Computer Security*.
12. M. Bugliesi, R. Focardi, and M. Maffei. Compositional analysis of authentication protocols. In *Proceedings of European Symposium on Programming (ESOP 2004)*, volume 2986 of *Lecture Notes in Computer Science*, pages 140–154. Springer-Verlag, 2004.
13. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL '77: Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 238–252, New York, NY, USA, 1977. ACM Press.
14. F. Crazzolaro and G. Winskel. Events in security protocols. In *CCS '01: Proceedings of the 8th ACM conference on Computer and Communications Security*, pages 96–105, New York, NY, USA, 2001. ACM Press.
15. D. Denning and G. Sacco. Timestamps in key distributed protocols. *Communication of the ACM*, 24(8):533–535, 1981.
16. Shaddin F. Doghmi, Joshua D. Guttman, and F. Javier Thayer. Searching for shapes in cryptographic protocols, 2007. To appear in *Proceedings of Thirteenth International Conference on tools and algorithms for the construction and analysis of systems (TACAS 2007)*.
17. D. Dolev and A. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
18. Laboratoire Spécification 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 *Proceedings of European Symposium on Programming (ESOP 2005)*, LNCS, pages 141–156. Springer-Verlag, 2005.
22. A. D. Gordon and A. Jeffrey. Types and effects for asymmetric cryptographic protocols. *Journal of Computer Security*, 12(3/4):435–484, 2004.
23. A. D. Gordon and A. Jeffrey. Secrecy despite compromise: Types, cryptography, and the pi-calculus. In *Proceedings of CONCUR 2005 - Concurrency Theory, 16th International Conference*, volume 3653, pages 186–201. Springer-Verlag, 2005.
24. J. D. Guttman and F. J. Thayer. Authentication tests and the structure of bundles. *Theoretical Computer Science*, 283(2):333–380, 2002.
25. J. D. Guttman, F.J. Thayer, J. A. Carlson, J. C. Herzog, J. D. Ramsdell, and B. T. Sniffen. Trust management in strand spaces: a rely-guarantee method. In *Proceedings of European Symposium on Programming (ESOP 2004)*, volume 2986 of *Lecture Notes in Computer Science*, pages 325–339. Springer-Verlag, 2004.
26. G. Lowe. “A Hierarchy of Authentication Specification”. In *Proceedings of the 10th Computer Security Foundation Workshop (CSFW'97)*, pages 31–44. IEEE Computer Society Press, 1997.
27. R. M. Needham and M. D. Schroeder. Authentication revisited. *ACM SIGOPS Operating Systems Review*, 21(1):7–7, 1987.

Table 6 Transition System for ρ -spi

Transition rules: We omit the symmetric rule of PAR.

<p style="text-align: center; margin: 0;">NAME RES</p> $\frac{n \notin \text{bn}(s) \cup \text{fn}(s)}{\langle s, \text{new}(n).P \rangle \rightarrow \langle s :: \text{new}(n), P \rangle}$	<p style="text-align: center; margin: 0;">KEY-PAIR RES</p> $\frac{k_I^-, k_I^+ \notin \text{bn}(s) \cup \text{fn}(s)}{\langle s, \text{new}^\mp(k_I).P \rangle \rightarrow \langle s :: \text{new}^\mp(k_I), P \rangle}$
<p style="text-align: center; margin: 0;">INPUT</p> $\frac{s \vdash G \quad \sigma = \text{bind}(M, G) \neq \uparrow}{\langle s, \text{in}(M).P \rangle \rightarrow \langle s :: \text{in}(M\sigma), P\sigma \rangle}$	<p style="text-align: center; margin: 0;">OUTPUT</p> $\langle s, \text{out}(G).P \rangle \rightarrow \langle s :: \text{out}(G), P \rangle$
<p style="text-align: center; margin: 0;">BEGIN</p> $\langle s, \text{begin}_G^i(A, I, G_1; G_2).P \rangle \rightarrow \langle s :: \text{begin}_G^i(A, I, G_1; G_2), P \rangle$	
<p style="text-align: center; margin: 0;">END</p> $\langle s, \text{end}_G^i(A, I, G_1; G_2).P \rangle \rightarrow \langle s :: \text{end}_G^i(A, I, G_1; G_2), P \rangle$	<p style="text-align: center; margin: 0;">PRINCIPAL</p> $\frac{\langle s, P \rangle \rightarrow \langle s :: t, P' \rangle}{\langle s, A \triangleright P \rangle \rightarrow \langle s :: t, A \triangleright P' \rangle}$
<p style="text-align: center; margin: 0;">PAR</p> $\frac{\langle s, P \rangle \rightarrow \langle s', P' \rangle}{\langle s, P Q \rangle \rightarrow \langle s', P' Q \rangle}$	<p style="text-align: center; margin: 0;">REPLICATION</p> $\langle s, !P \rangle \rightarrow \langle s, P \mid !P \rangle$

28. B. Clifford Neuman and Theodore Ts'o. Kerberos : An authentication service for computer networks. *IEEE Communication*, 32(9):33–38, 1994.
29. J. L. Peterson. Petri nets. *ACM Comput. Surv.*, 9(3):223–252, 1977.
30. M. D. Schroeder R. M. Needham. Using encryption for authentication in large networks of computers. *ACM Communication*, 21(12):993–999, 1978.
31. D. Wagner and B. Schneier. Analysis of the SSL 3.0 protocol. In *Proceedings of 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, New York, NY, USA, 1987. Springer-Verlag New York, Inc.
33. T.Y.C. Woo and S.S. Lam. “A Semantic Model for Authentication Protocols”. In *Proceedings of 1993 IEEE Symposium on Security and Privacy*, pages 178–194, 1993.

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 $s \in \text{Act}^*$ is a trace and P is a closed process. Each transition $\langle s, P \rangle \rightarrow \langle s :: t, P' \rangle$ simulates one computation step in P and records the corresponding action in the trace. The set of all possible actions, noted *Act*, includes the action $\text{new}(n)$ generated by name restriction and $\text{new}^\mp(k_I)$ by key-pair restriction, $\text{in}(G)$ generated by input, $\text{out}(G)$ by output, $\text{begin}_G^i(A, I, G_1; G_2)$ and $\text{end}_G^i(A, I, G_1; G_2)$ by ‘begin’ and ‘end’, respectively. 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 $[G_1/x_1, \dots, G_n/x_n]$.

Table 7 Deductive system and binding

Message Manipulation Rules			
OUT $\frac{out(G) \in s}{s \vdash G}$	ENV $\frac{a \notin \text{bn}(s)}{s \vdash a}$	PAIR $\frac{s \vdash G_1 \quad s \vdash G_2}{s \vdash (G_1, G_2)}$	PAIR DES $\frac{s \vdash (G_1, G_2)}{s \vdash G_i, i \in [1, 2]}$
ENCRYPTION $\frac{s \vdash G \quad s \vdash a}{s \vdash \{G\}_a}$		DECRYPTION $\frac{s \vdash \{G\}_a \quad s \vdash \bar{a}}{s \vdash G}$	
PUBLIC KEYS $s \vdash k_I^+$	ENEMY KEYS $s \vdash k_{EI}$	ENEMY KEYS $s \vdash k_{IE}$	ENEMY KEYS $s \vdash k_E^-$

Binding

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

The 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 $\bar{\bar{n}} = n$, $\overline{k_{IJ}} = k_{IJ}$, $\bar{I} = I$, $k_I^+ = k_I^-$ and $\overline{k_I^-} = k_I^+$. OUTPUT, BEGIN and END are self-explanatory. Finally, PRINCIPAL adds the principal name to the performed action, 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 to the different notation for encryption. 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 either an atomic name a ; 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 deductive 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 may construct and destruct pairs. By ENCRYPTION, and DECRYPTION the environment can encrypt and decrypt messages only knowing the required keys. By PUBLIC KEYS, all the public keys are known by 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 start authentication sessions and, generally speaking, to interact with the other participants by pretending to be a trusted principal.

A.1 Process well-formedness

We remark that the ρ -spi syntax is quite liberal, giving the possibility of writing nonsense processes like $A \triangleright B \triangleright P$, where an identity declaration nests within another one, or $A \triangleright \text{begin}_n(B, I, M)$, where A asserts a begin-event in place of B . Furthermore, it makes sense to require that a shared key k_{IJ} should only be used by I and J and a private key k_I^- should only be used by I . Although this is not required by the analysis, the ρ -spi calculus comes with a notion of process well-formedness ruling out the above mentioned undesired process behaviors and enforcing the correct use of long-term keys. It is formalized through the standard notion of scope: we say that the *scope* of $A \triangleright$ in process $A \triangleright P$, is process P .

Definition 10 (Process Well-Formedness). *A process P is well-formed if the following conditions hold:*

1. *All the occurrences of $A \triangleright$ in P are not in the scope of some other $B \triangleright$;*
2. *Every $\text{begin}_N^i(A, I, M)$ in P is in the scope of $A \triangleright$;*
3. *Every $\text{end}_N^i(B, J, M)$ in P is in the scope of $B \triangleright$;*
4. *Every symmetric key k_{IJ} occurring in P is in the scope of either $I \triangleright$ or $J \triangleright$;*
5. *Every private key k_I^- occurring in P is in the scope of $I \triangleright$.*

Well-formedness is trivially verifiable by a simple syntactic inspection.

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 deductive system similar to the one of ρ -spi calculus and reported in Table 8. The environment knows every term sent on the network (OUT) and the special name \mathcal{E} (ENV MSG), it can construct and destruct pairs (PAIR and PAIR DES) and encrypt and decrypt terms only knowing the required keys (ENCRYPTION and DECRYPTION). 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_{abs} , 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,

Table 8 Deductive System and Binding

Notation: In OUT, PAIR DES, DECRYPTION, $\#1$ s.t. $\lceil G \rceil \in \{1, k_1^+, k_E^-, k_{IE}, k_{EI}, \mathcal{E}\}$.
 \bar{v} is defined as $\bar{n} = n$, $\bar{1} = 1$, $\bar{\mathcal{E}} = \mathcal{E}$, $k_1^{\bar{\dagger}} = k_1^-$ and $k_1^- = k_1^{\dagger}$.

$\frac{\text{OUT}}{\frac{G \in \mathcal{G}}{\mathcal{G} \vdash G}}$	$\frac{\text{ENV MSG}}{\mathcal{G} \vdash \mathcal{E}}$	$\frac{\text{PAIR}}{\frac{\mathcal{G} \vdash G_1 \quad \mathcal{G} \vdash G_2}{\mathcal{G} \vdash (G_1, G_2)}}$	$\frac{\text{PAIR DES}}{\frac{\mathcal{G} \vdash (G_1, G_2)}{\mathcal{G} \vdash G_i \quad \forall i \in [1, 2]}}$
$\frac{\text{ENCRYPTION}}{\frac{\mathcal{G} \vdash G \quad \mathcal{G} \vdash v}{\mathcal{G} \vdash \{G\}_v}}$	$\frac{\text{DECRYPTION}}{\frac{\mathcal{G} \vdash \{G\}_v \quad \mathcal{G} \vdash v' \quad \text{bind}_{\text{abs}}(v', \bar{v}) \neq \uparrow \vee \text{bind}_{\text{abs}}(\bar{v}, v) \neq \uparrow}{\mathcal{G} \vdash G}}$		

Binding

$\text{bind}_{\text{abs}}(v_1, v_2) = v_2$	if either $[v_1] = [v_2]$ or $([v_1] = \mathcal{E} \wedge [v_2] \in \{1, k_1^+, k_{IE}, k_{EI}, k_E^-, \mathcal{E}\})$ or $([v_2] = \mathcal{E} \wedge [v_1] \in \{1, k_1^+, k_{IE}, k_{EI}, k_E^-, \mathcal{E}\})$
$\text{bind}_{\text{abs}}(x, v) = v_{(x)}$	
$\text{bind}_{\text{abs}}((M, M'), (G, G')) = (G_1, G_2)$	if $\text{bind}_{\text{abs}}(M, G) = G_1 \wedge \text{bind}_{\text{abs}}(M', G') = G_2$
$\text{bind}_{\text{abs}}(\{M\}_v, \{G\}_{v'}) = \{G'\}_{v'}$	if $\text{bind}_{\text{abs}}(M, G) = G' \wedge$ $(\text{bind}_{\text{abs}}(\bar{v}, v') \neq \uparrow \vee \text{bind}_{\text{abs}}(v, \bar{v}') \neq \uparrow)$
$\text{bind}_{\text{abs}}(M, G) = \uparrow$	otherwise

yields the the term G' obtained by labelling G according to the variables in M . If pattern-matching fails, bind_{abs} returns \uparrow . Function bind_{abs} 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_{abs} 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 .

C Proofs of Soundness and Safety

This section proves the main results of our static analysis technique. For proving soundness results we need to refine the concretization function for terms. In particular, the definition in Table 4 says that the occurrences of the same name in an abstract term are instantiated by the same ρ -spi name. This definition can be soundly used for the concretization of paths and causal graphs into traces and configurations, respectively, but it is too precise for the concretization of terms in the knowledge of the environment. For instance, if the abstract environment knows (n^j, n^j) , where j is an index, then this means that the abstract name n is leaked out. In the ρ -spi calculus, this means that the environment may combine different protocol sessions, thus knowing terms like (n_1, n_2) , where $\{n_1, n_2\} \subseteq \gamma_{\text{trm}}(n)$ and the same abstract name is concretized into two different ρ -spi names. For tackling this problem, we relax the concretization of abstract (ground)

Table 9 Result of $paths(graph(Prot), end_n^1(A, B, m), \emptyset)$

$$\left. \begin{array}{l} \overbrace{new^{\bar{+}}(k_A) :: new(m) :: new(n)}^a :: out(\{B, n, m\}_{k_A^+}) :: in(n) \\ \overbrace{new^{\bar{+}}(k_A) :: in(\{B, n(x), m(z)\}_{k_A^-}) :: begin_{n(x)}^1(A, B, m(z)) :: out_{com}(n(x), n^{i5})}^b :: in(n) \\ \overbrace{a :: out_{com}(\{B, n, m\}_{k_A^+}, \{B^{i1}, n_{(x)}^{i2}, m_{(z)}^{i3}\}_{k_A^+ i_A}) :: in(\{B, n(x), m(z)\}_{k_A^-}) :: begin_{n(x)}^1(A, B, m(z)) :: out_{com}(n(x), n^{i5})}^c :: in(n) \end{array} \right\}$$

$$\left. \begin{array}{l} a :: out(\{B, n, m\}_{k_A^+}) :: in(n) \\ \overbrace{new^{\bar{+}}(k_A) :: in(\{B, n(x), \mathcal{E}(z)\}_{k_A^-}) :: begin_{n(x)}^1(A, B, \mathcal{E}(z)) :: out_{com}(n(x), n^{i5})}^d :: in(n) \\ \overbrace{b :: out_{com}(n(x), n^{i5})}^e :: in(\{B, n(x), \mathcal{E}(z)\}_{k_A^-}) :: begin_{n(x)}^1(A, B, \mathcal{E}(z)) :: out_{com}(n(x), n^{i5})}^e :: in(n) \\ c :: e :: out_{com}(n(x), n^{i5}) :: in(n) \end{array} \right\}$$

$$\left. \begin{array}{l} a :: out(\{B, n, m\}_{k_A^+}) :: in(n) \\ d :: out_{com}(n(x), n^{i5}) :: in(n) \\ \overbrace{new^{\bar{+}}(k_A) :: in(\{\mathcal{E}, n(x), n(z)\}_{k_A^-}) :: begin_{n(x)}^1(A, B, n(z))}^f :: e :: out_{com}(n(x), n^{i5})}^f :: in(n) \\ \overbrace{b :: out_{com}(n(x), n^{i5})}^g :: in(\{B, n(x), n(z)\}_{k_A^-}) :: begin_{n(x)}^1(A, B, n(z))}^g :: e :: out_{com}(n(x), n^{i5})}^g :: in(n) \\ \overbrace{b :: out_{com}(n(x), n^{i5})}^i :: in(\{B, n(x), n(z)\}_{k_A^-}) :: begin_{n(x)}^1(A, B, n(z))}^i :: e :: out_{com}(n(x), n^{i5})}^i :: in(n) \\ c :: g :: e :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: i :: e :: out_{com}(n(x), n^{i5}) :: in(n) \end{array} \right\}$$

$$\left. \begin{array}{l} a :: out(\{B, n, m\}_{k_A^+}) :: in(n) \\ f :: out_{com}(n(x), n^{i5}) :: in(n) \\ b :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ b :: i :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: i :: out_{com}(n(x), n^{i5}) :: in(n) \end{array} \right\} \quad \left. \begin{array}{l} a :: out(\{B, n, m\}_{k_A^+}) :: in(n) \\ f :: out_{com}(n(x), n^{i5}) :: in(n) \\ d :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ d :: i :: out_{com}(n(x), n^{i5}) :: in(n) \\ b :: e :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ b :: e :: i :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: e :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: e :: i :: out_{com}(n(x), n^{i5}) :: in(n) \end{array} \right\}$$

$$\left. \begin{array}{l} a :: out(\{B, n, m\}_{k_A^+}) :: in(n) \\ f :: out_{com}(n(x), n^{i5}) :: in(n) \\ d :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ d :: i :: out_{com}(n(x), n^{i5}) :: in(n) \\ b :: i :: out_{com}(n(x), n^{i5}) :: in(n) \\ b :: e :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: i :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: e :: g :: out_{com}(n(x), n^{i5}) :: in(n) \end{array} \right\} \quad \left. \begin{array}{l} a :: out(\{B, n, m\}_{k_A^+}) :: in(n) \\ f :: out_{com}(n(x), n^{i5}) :: in(n) \\ d :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ d :: i :: out_{com}(n(x), n^{i5}) :: in(n) \\ b :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ b :: e :: i :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: g :: out_{com}(n(x), n^{i5}) :: in(n) \\ c :: e :: i :: out_{com}(n(x), n^{i5}) :: in(n) \end{array} \right\}$$

terms as follows so as to concretize different occurrences of the same abstract name into the same ρ -spi name only when they occur within an integer component, which cannot be forged by the environment. Function $\gamma_{\text{env}} : \mathcal{C}, G \mapsto \{G_1, \dots, G_n\}$ is defined as follows:

$$\gamma_{\text{env}}(\mathcal{C}, G) \triangleq \begin{cases} \gamma_{\text{trm}}(v) & \text{if } G = v \\ \{(G_1, G_2) \mid G_1 \in \gamma_{\text{env}}(\mathcal{C}, G_1) \wedge G_2 \in \gamma_{\text{env}}(\mathcal{C}, G_2)\} & \text{if } G = (G_1, G_2) \\ \{\{G\}_v \mid G \in \gamma_{\text{env}}(\mathcal{C}, G') \wedge v \in \gamma_{\text{env}}(\mathcal{C}, v)\} & \text{if } G = \{G'\}_v \\ & \wedge \text{output}(\text{nodes}(\mathcal{C})) \vdash v \\ \gamma_{\text{trm}}(G) & \text{otherwise} \end{cases}$$

The following lemma says that $\gamma_{\text{env}}(\mathcal{C}, G)$ contains $\gamma_{\text{trm}}(G)$.

Lemma 1 (Term Concretization). *For every term G and causal graph \mathcal{C} , $\gamma_{\text{trm}}(G) \subseteq \gamma_{\text{env}}(\mathcal{C}, G)$.*

Proof. We reason by induction on the structure of G .

Base Case The base case is $G = v$, which is trivial since, by definition of γ_{env} , we have that $\gamma_{\text{env}}(\mathcal{C}, v) = \gamma_{\text{trm}}(v)$ for any \mathcal{C} .

Inductive Step Let us suppose that $G = (G_1, G_2)$, for some G_1, G_2 . By definition of γ_{trm} , for every $(G_1, G_2) \in \gamma_{\text{trm}}(G)$, there exists σ such that $(G_1, G_2) = (G_1, G_2)\sigma$ and thus $G_1 = G_1\sigma$ and $G_2 = G_2\sigma$. Therefore, we have that $G_1 \in \gamma_{\text{trm}}(G_1)$ and $G_2 \in \gamma_{\text{trm}}(G_2)$. By induction hypothesis, $G_1 \in \gamma_{\text{env}}(\mathcal{C}, G_1)$ and $G_2 \in \gamma_{\text{trm}}(\mathcal{C}, G_2)$ for any \mathcal{C} and, by definition of γ_{env} , $(G_1, G_2) \in \gamma_{\text{env}}(\mathcal{C}, G)$, thus getting the result.

The proof for pairs and ciphertexts $\{G'\}_v$, where $\text{output}(\text{nodes}(\mathcal{C})) \vdash v$, follows the same reasoning. If $G = \{G'\}_v$ and $\text{output}(\text{nodes}(\mathcal{C})) \not\vdash v$ then, by definition of γ_{env} , $\gamma_{\text{env}}(\mathcal{C}, G) = \gamma_{\text{trm}}(G)$ for any \mathcal{C} , as desired.

The next lemma says that if there exists a binding between two ρ -spi terms, then there exists a binding between their abstractions.

Lemma 2 (Binding). *Let \mathcal{C} be a causal graph, G and M ρ -spi terms, G and M abstract terms such that $G \in \gamma_{\text{env}}(\mathcal{C}, G)$, $M \in \gamma_{\text{trm}}(M)$ and $\text{bind}(M, G) \neq \uparrow$. Then $\exists G'$ such that $\text{bind}_{\text{abs}}(M, G) = G'$, $\sigma = \text{subst}(G')$ and $\overline{G} \in \gamma_{\text{trm}}(M\sigma)$.*

Proof. By induction on the structure of M .

Base Case We have two cases, depending on whether M is a variable or a name:

- Let $M = x$ and $G = v$. By definition of γ_{trm} and γ_{env} , $M = x$ and $G = v$, for some x and v such that $x \in \gamma_{\text{trm}}(x)$ and $v \in \gamma_{\text{env}}(\mathcal{C}, v) = \gamma_{\text{trm}}(v)$. By definition of bind_{abs} and subst , $\text{bind}_{\text{abs}}(x, v) = v_{(x)}$ and $\text{subst}(v_{(x)}) = [\overline{v}]_{(x)}/x$. By definition of γ_{trm} , $v \in \gamma_{\text{trm}}([\overline{v}]_{(x)})$, as desired.
- Let $M = a$ and $G = a$. By definition of γ_{trm} , $M = v$, for some v such that $a \in \gamma_{\text{trm}}(v)$. Similarly, by definition of γ_{env} , $G = v'$, for some v' such that $a \in \gamma_{\text{env}}(\mathcal{C}, v') = \gamma_{\text{trm}}(v')$. If $a \notin \mathcal{M}_{\mathcal{E}} \cup I\mathcal{D} \cup \mathcal{K}^+ \cup \mathcal{K}_{\mathcal{E}}$, then $[v] = [v']$ and, by definition of bind_{abs} , $\text{bind}_{\text{abs}}(v, v') = v'$, as desired. If $a \in \mathcal{M}_{\mathcal{E}} \cup I\mathcal{D} \cup \mathcal{K}^+ \cup \mathcal{K}_{\mathcal{E}}$, then v or v' may be the message \mathcal{E} and the reasoning is similar.

Inductive step We have two cases, depending on whether M is a pair or a ciphertext:

- Let $M = (M_1, M_2)$ and $G = (G_1, G_2)$. By definition of γ_{trm} , there exist M_1, M_2, σ' such that $M = (M_1, M_2)$ and $M\sigma' = M$. By definition of γ_{env} , there exist G_1, G_2 such that $G = (G_1, G_2)$, $G_1 \in \gamma_{\text{env}}(\mathcal{C}, G_1)$ and $G_2 \in \gamma_{\text{env}}(\mathcal{C}, G_2)$. By definition of function bind , $\text{bind}(M, G) = \text{bind}(M_1, G_1) \uplus \text{bind}(M_2, G_2)$. By induction hypothesis there exist $G'_1 = \text{bind}_{\text{abs}}(M_1, G_1)$ and $G'_2 = \text{bind}_{\text{abs}}(M_2, G_2)$ such that $\overline{G}_1 \in \gamma_{\text{trm}}(M_1\sigma_1)$ and $\overline{G}_2 \in \gamma_{\text{trm}}(M_2\sigma_2)$, where $\sigma_1 = \text{subst}(G'_1)$ and $\sigma_2 = \text{subst}(G'_2)$. Since $M\sigma' = M$ and $\text{bind}(M, G) \neq \uparrow$, it turns out that for every $v \in \text{dom}(\sigma_1) \cap \text{dom}(\sigma_2)$, $\sigma_1(v) = \sigma_2(v)$. Thus $\sigma = \sigma_1 \uplus \sigma_2 \neq \uparrow$ and $(G_1, G_2) \in \gamma_{\text{trm}}((M_1, M_2)\sigma)$, as desired.
- Let $M = \{M'\}_{a_1}$ and $G = \{G'\}_{a_2}$, with $\overline{a_1} = a_2$. By definition of γ_{trm} , there exists M', v_1, σ' such that $M = \{M'\}_{v_1}$ and $M\sigma' = M$. By definition of γ_{env} , there exists G', v_2 such that $G = \{G'\}_{v_2}$. We have two cases: either $\text{output}(\text{nodes}(\mathcal{C})) \vdash v_2$, and thus $\{G'\}_{a_2} \in \gamma_{\text{trm}}(\{G'\}_{v_2})$, or $\text{output}(\text{nodes}(\mathcal{C})) \not\vdash v_2$, and thus $G' \in \gamma_{\text{env}}(G')$ and $a_2 \in \gamma_{\text{trm}}(v_2)$. In the former case the result is immediate. In the latter case, the proof is similar to the one for pairs. Notice that if $\overline{a_1} = a_2$ then either $\text{bind}_{\text{abs}}(\overline{v_1}, v_2) \neq \uparrow$ or $\text{bind}_{\text{abs}}(v_1, \overline{v_2}) \neq \uparrow$.

The following lemma states that causal graphs are a sound abstraction of ρ -spi configurations as far as secrecy is concerned. In particular, if a message G is known to the environment associated with the ρ -spi configuration $\langle s, P \rangle$, namely $s \vdash G$, then an abstraction G of G is known to the environment in the causal graph abstracting $\langle s, P \rangle$. Furthermore, every integer component in G is associated with an output in s and the instantiation of messages occurring therein is consistent.

Lemma 3 (Abstract Environment). *Suppose $\langle s, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$ and $s \vdash G$. Then $\exists G$ such that the following statements hold:*

1. $\text{output}(\text{nodes}(\mathcal{C})) \vdash G$ and $G \in \gamma_{\text{env}}(\mathcal{C}, G)$
2. if $\text{bind}(G, M) \neq \uparrow$ and $M \in \gamma_{\text{trm}}(M)$, then $\exists G_{\text{lab}}$ such that $\text{bind}_{\text{abs}}(G, M) = G_{\text{lab}}$, $\sigma = \text{subst}(G_{\text{lab}})$ and $\overline{G} \in \gamma_{\text{trm}}(M\sigma)$
3. $\forall G' \in \text{int}(\mathcal{C}, G_{\text{lab}})$, $\exists \text{out}(G'').P'' \in \text{nodes}(\mathcal{C})$ and G'' such that
 - (a) $[G'] \in \text{terms}(\downarrow G'')$
 - (b) $s = t_1 :: \dots :: t_{r-1} :: \text{out}(G'') :: t_{r+1} :: \dots :: t_n$
 - (c) $\forall i \in [1, n]$, $\exists Q_i \in \text{nodes}(\mathcal{C})$, $\mathcal{S}_i \in \text{paths}(\mathcal{C}, Q_i, \emptyset)$, $(\mathcal{S}_i, \sigma_i) \in \gamma_{\text{pset}}(\mathcal{S}_i)$ s.t. $\mathcal{S}_i = \mathcal{S}_j :: t_i$, with $j < i$
 - (d) $Q_r \in \text{threads}(P'')$
 - (e) Let σ_1 be the least substitution such that $\overline{G} = (M\sigma)\sigma_1$. If $\sigma_1(v) = n$ and $v^{\text{out}(G'').P.k} \in \text{terms}(G')$, then $\sigma_r(\pi^k(G'')) = n$

Proof. By induction on the derivation of $s \vdash G$.

Base Case We have four cases depending on the derivation rule applied (OUT, PUBLIC KEYS, ENEMY KEYS or ENV):

OUT The judgement $s \vdash G$ is proved by OUT and $s = t_1 :: \dots :: t_{r-1} :: \text{out}(G) :: \dots :: t_{r+1} :: \dots :: t_n$. Since $\langle s, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, the first condition in the definition of γ_{net} proves item 3c. By definition of function paths , we have that $\text{out}(G).P \in$

$nodes(\mathcal{C})$ and $Q_r \in threads(P)$, with $G \in \gamma_{trm}(G)$. Therefore $output(nodes(\mathcal{C})) \vdash G$ and, by Lemma 1, $G \in \gamma_{env}(\mathcal{C}, G)$, thus proving item 1. By Lemma 2, item 2 holds as well. Finally, item 3 is trivially satisfied since all integer components in G are output by the same process $out(G).P$.

PUBLIC KEYS, ENEMY KEYS, ENV By ENV MSG $output(nodes(\mathcal{C})) \vdash \mathcal{E}$ and $a \in \gamma_{env}(\mathcal{C}, \mathcal{E})$, whenever a is a public key, an enemy key, an identity or a message generated by the environment. Item 2 is proved by Lemma 2. The proof of the third item is trivial since, for any v such that $[v] = \mathcal{E}$, $int(\mathcal{C}, v) = \emptyset$.

Inductive Step We have four cases according to the last derivation rule applied (PAIR, PAIR DES, DECRYPTION or ENCRYPTION):

PAIR, PAIR DES, DECRYPTION The thesis is straightforwardly proved by induction hypothesis.

ENCRYPTION Let us suppose that $G = \{G'\}_{v'}$, with $s \vdash G'$ and $s \vdash v'$. By induction hypothesis, there exist G' and v' such that $G' \in \gamma_{env}(\mathcal{C}, G')$, $v' \in \gamma_{env}(\mathcal{C}, v')$, $output(nodes(\mathcal{C})) \vdash G'$ and $output(nodes(\mathcal{C})) \vdash v'$. By ENCRYPTION, we have that $output(nodes(\mathcal{C})) \vdash \{G'\}_{v'}$ and, by definition of γ_{env} , we get $\{G'\}_{v'} \in \gamma_{env}(\mathcal{C}, \{G'\}_{v'})$, thus proving item 1. Item 2 is proved by Lemma 2. Let $G_{lab} = \{G_1\}_{v_1}$: since $output(nodes(\mathcal{C})) \vdash v_1$, $int(\mathcal{C}, \{G_1\}_{v_1}) = int(\mathcal{C}, G_1) \cup int(\mathcal{C}, v_1)$ and item 3 is trivially proved by induction hypothesis.

The next definition introduces the notion of path set approximation, which is used in Lemma 4 for proving that cycles may be soundly abstracted away in cycle-invariant causal graphs.

Definition 11 (Path set approximation). We say that S' is an approximation of S , written $S \preceq S'$, if and only if

1. for every $s' :: out(G_1) :: in_{com}(G_2, G) \in S'$ there exists $s \in S$ such that:
 - either $s = s'$
 - or $s = s' :: out(G_1) :: in_{com}(G_3, G) :: s''$, for some s'' and G_3 , and $s' \vdash (v, j) = (v', j')$ and $v \in msgs(G_1)$ imply $s \vdash (v, i) = (v', j')$, with $|s| = i$ and $|s'| = j$.
2. for every $s' :: in(G) \in S'$, there exists $s \in S$ such that:
 - either $s = s'$
 - or $s = s' :: in(G) :: s''$, for some s'' , and $s = s' :: out(G_1) :: in_{com}(G_3, G) :: s''$, for some s'' , and $s' \vdash (v, j) = (v', j')$ implies $s \vdash (v, i) = (v', j')$, with $|s| = i$ and $|s'| = j$.

Lemma 4 (Cycle-invariant causal graphs). Let \mathcal{C} be a cycle-invariant causal graph and $out(G_1).P_1 \xrightarrow{G_1} \boxed{com(G)} \rightarrow Q$ and $in(M).P \xrightarrow{} \boxed{com(G)} \rightarrow Q$ be edges in \mathcal{C} .

For every $S \in paths(\mathcal{C}, out(G_1).P, \emptyset)$ there exists $S' \in paths(\mathcal{C}, Q, \emptyset)$ such that $S' \preceq S$.
For every $S \in paths(\mathcal{C}, in(M).P, \emptyset)$ there exists $S' \in paths(\mathcal{C}, Q, \emptyset)$ such that $S' \preceq S$.

Proof. We prove the first item as the proof for the second one is similar. Let us suppose that there exists a cycle in the causal graph containing the edge $in(M).P \xrightarrow{} \boxed{com(G)} \rightarrow Q$. Otherwise, the thesis is trivially satisfied. More formally, there exists $S \in paths(\mathcal{C},$

$\text{out}(G_1).P, \emptyset$) such that $\mathcal{S}' = \{s \in \mathcal{S} \mid \exists s_1, s_2 \text{ s.t. } s = s_1 :: \text{out}(G_1) :: \text{in}_{\text{com}}(G'_1, G) :: s_2\} \neq \emptyset$.

By definition of function *paths*, $s_1 :: \text{out}(G_1) :: \text{in}_{\text{com}}(G'_1, G) :: s_2 \in \mathcal{S}'$ if and only if there exist $\mathcal{E}, \mathcal{S}'' \in \text{paths}(\mathcal{C}, Q, \mathcal{E})$ such that $s_1 :: \text{out}(G_1) :: \text{in}_{\text{com}}(G'_1, G) \in \mathcal{S}''$. By definition of function *paths*, $\mathcal{S}'' \in \text{paths}(\mathcal{C}, Q, \emptyset)$. Since \mathcal{G} is cycle-invariant, there exists $s' = s_1 :: \text{out}(G_1, G'_1) :: \text{in}(G) :: s'_2 \in \mathcal{S}'$ such that $s' \vdash (v, j) = (v', j')$ and $v \in \text{msgs}(G_1)$ imply $s \vdash (v, i) = (v', j')$, with $|s| = i$ and $|s'| = j$, thus proving the thesis.

The next proposition states that if a configuration is an instance of a causal graph guaranteeing authenticity, the trace in the configuration guarantees weak authenticity. In the following, we write $t_1 :: \dots :: t_n \sqsubseteq s$ to say that there exist s_1, \dots, s_{n+1} such that $s = s_1 :: t_1 :: \dots :: s_n :: t_n :: s_{n+1}$.

Lemma 5 (Authenticity). *If $\langle s, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$ and \mathcal{C} guarantees authenticity, then s guarantees weak authenticity.*

Proof. By condition (i) in the definition of γ_{net} , if $\exists i$ such that $s = s_1 :: \text{end}_{G_1}^i(A, B, G_2) :: s_2$ then there exist $Q \in \text{nodes}(\mathcal{C})$, $\mathcal{S} \in \text{paths}(\mathcal{C}, Q, \emptyset)$ such that, for every $s \in \mathcal{S}$, there exists $s' \sqsubseteq s_1$ such that $s' :: \text{end}_N^i(A, B, M) \in \gamma_{\text{path}}(s)$. Since the causal graph guarantees authenticity, there exist G_1, G_2 such that $s' :: \text{end}_{G_1}^i(A, B, G_2) \in \mathcal{S}$ guarantees authenticity. Thus there exist G'_1, G'_2, s'_1, s'_2 such that $s' = s'_1 :: \text{begin}_{G'_1}^i(A, B, G'_2) :: s'_2$, $G_1 \simeq_{s, i, j} G'_1$ and $G_2 \simeq_{s, i, j} G'_2$, where $|s'_1 :: \text{begin}_{G'_1}^i(A, B, G'_2)| = j$ and $|s' :: \text{end}_{G_1}^i(A, B, G_2)| = i$. By definition of γ_{path} , the abstract messages occurring in the begin and end assertions have the same instantiation and thus $\text{begin}_N^i(A, B, M) \in s'$ as desired.

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

Proof. Trivial, by definition of γ_{net} and by observing that the trace in the configuration is empty.

The preservation theorem states that the set of configurations abstracted by a causal graph is closed under process reduction. This means that causal graphs are a sound model for statically reasoning on security properties of ρ -spi processes.

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

Proof. The proof proceeds by induction on the length of the derivation for $\langle s, P \rangle \rightarrow \langle s', P' \rangle$: if P is a parallel composition or a process with identifier this may take several steps. The base cases are given by the remaining transition rules. Since $\langle s, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, with $s = t_1 :: \dots :: t_n$, we have that

1. $\forall i \in [1, n], \exists Q_i \in \text{nodes}(\mathcal{C}), \mathcal{S}_i \in \text{paths}(\mathcal{C}, Q_i, \emptyset), (S_i, \sigma_i) \in \gamma_{\text{pset}}(S_i) \text{ s.t. } S_i = S_j :: t_i, \text{ with } j < i$
2. $\forall Q \in \text{threads}(P), \exists i \in [1, n], \sigma \text{ s.t. } Q = Q_i \sigma \wedge \sigma \uplus \sigma_i \neq \uparrow$

New $\langle s, \text{new}(n).P \rangle \rightarrow \langle s :: \text{new}(n), P \rangle$

For proving $\langle s :: \text{new}(n), P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, we have to prove the two conditions in the definition of γ_{net} .

Condition (i) Since $\langle \text{new}(n).P, s \rangle \in \gamma_{\text{net}}(\mathcal{C})$, there exist $i \in [1, n]$ and $Q_i = \text{new}(n).P \in \text{nodes}(\mathcal{C})$ such that $\text{new}(n).P = \text{new}(n).P\sigma$, for some σ such that $\sigma \uplus \sigma_i \neq \uparrow$. By definition of causal graph, if $\text{new}(n).P \in \text{nodes}(\mathcal{C})$, then there exists $Q \in \text{nodes}(\mathcal{C})$ such that $Q \in \text{threads}(P)$ and $\text{new}(n).P \rightarrow Q \in \text{edges}(\mathcal{C})$. By definition of function *paths*, $S_i :: \text{new}(n) \in \text{paths}(\mathcal{C}, Q, \emptyset)$. By function γ_{pset} , $(S_i :: \text{new}(n), \sigma_i \cup [n/n]) \in \gamma_{\text{pset}}(S_i :: \text{new}(n))$, thus proving condition (i).

Condition (ii) Since $\text{new}(n).P = \text{new}(n).P\sigma$, we get $P = P\sigma$ and $\sigma \uplus (\sigma_i \cup [n/n]) \neq \uparrow$, as desired.

Thus $\langle s :: \text{new}(n), P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, as desired. The reasoning for key-pair restrictions is similar.

Struct $\langle s, p.P \rangle \rightarrow \langle s :: p, P \rangle$,
with $p \in \{\text{out}(G), \text{begin}_{G_1}^i(A, I, G_2), \text{end}_{G_1}^i(A, I, G_2)\}$

As before, we prove distinctly the two conditions for $\langle s :: p, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$.

Condition(i) Since $\langle s, p.P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, there exist $i \in [1, n]$ and $Q_i = p.P \in \text{nodes}(\mathcal{C})$ such that $p.P = p.P\sigma$, for some σ such that $\sigma \uplus \sigma_i \neq \uparrow$. By definition of causal graph, if $p.P \in \text{nodes}(\mathcal{C})$, then there exists $Q \in \text{nodes}(\mathcal{C})$ such that $Q \in \text{threads}(P)$ and $p.P \rightarrow Q \in \text{edges}(\mathcal{C})$. By definition of function *paths*, $S_i :: p \in \text{paths}(\mathcal{C}, Q, \emptyset)$. Since $p.P = p.P\sigma$, we have that $p = p\sigma$: by definition of function γ_{pset} , $(S_i :: p, \sigma \cup \sigma_i) \in \gamma_{\text{pset}}(S_i :: p)$, thus proving condition (i).

Condition (ii) Since $p.P = p.P\sigma$, we have that $P = P\sigma$, as desired.

Thus $\langle s :: p, P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, as desired.

In $\langle s, \text{in}(M).P \rangle \rightarrow \langle s :: \text{in}(M\sigma), P\sigma \rangle$, with $\sigma = \text{bind}(M, G)$

Since $\langle s, \text{in}(M).P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, by condition (ii) there exist M and P such that $\text{in}(M).P \in \text{nodes}(\mathcal{C})$ and $M \in \gamma_{\text{trm}}(M)$. In the following, we reason on an arbitrary $Q \in \text{threads}(P\sigma)$. By Lemma 3, if $s \vdash G$, then there exists a term G such that

1. $\text{output}(\text{nodes}(\mathcal{C})) \vdash G$ and $G \in \gamma_{\text{env}}(\mathcal{C}, G)$
2. $\exists \sigma, G_{\text{lab}}$ such that $\text{bind}_{\text{abs}}(G, M) = G_{\text{lab}}$, $\sigma' = \text{subst}(G_{\text{lab}})$ and $\overline{G} \in \gamma_{\text{trm}}(M\sigma')$
3. $\forall G' \in \text{int}(\mathcal{C}, G_{\text{lab}})$, $\exists \text{out}(G'').P'' \in \text{nodes}(\mathcal{C})$, G'' and S such that
 - (a) $\lfloor G' \rfloor \in \text{terms}(\lfloor G'' \rfloor)$
 - (b) $s = t_1 :: \dots :: t_{r-1} :: \text{out}(G'') :: t_{r+1} :: \dots :: t_n$
 - (c) $\forall i \in [1, n]$, $\exists Q_i \in \text{nodes}(\mathcal{C})$, $S_i \in \text{paths}(\mathcal{C}, Q_i, \emptyset)$, $(S_i, \sigma_i) \in \gamma_{\text{pset}}(S_i)$ s.t. $S_i = S_j :: t_i$, with $j < i$
 - (d) $Q_r \in \text{threads}(P'')$
 - (e) Let σ_1 be the least substitution such that $\overline{G} = (M\sigma')\sigma_1$. If $\sigma_1(v) = n$ and $\sqrt{\text{out}(G'').P.i} \in \text{terms}(G')$, then $\sigma_r(\pi^j(G'')) = n$

By definition of causal graph, there exists $Q \in \text{nodes}(\mathcal{C})$ such that $Q \in \text{threads}(P\sigma')$, $\text{in}(M).P \xrightarrow{\text{com}(G)} Q \in \text{edges}(\mathcal{C})$. As before, we prove the two conditions for $\langle s :: \text{in}(M\sigma), P\sigma \rangle \in \gamma_{\text{net}}(\mathcal{C})$.

Condition (i) Since $\langle s, \text{in}(M).P \rangle \in \gamma_{\text{net}}(\mathcal{C})$, by condition (ii) applied on the source configuration, there exist $S_i \in \text{paths}(\mathcal{C}, \text{in}(M).P, \emptyset)$, S_i such that $S_i \in \gamma_{\text{pset}}(S_i)$ and $\text{in}(M).P = \text{in}(M).P\sigma''$, for some σ'' such that $\sigma'' \uplus \sigma_i \neq \uparrow$. By definition of γ_{pset} , $(S_i :: \text{in}(M\sigma), \sigma_1 \uplus \sigma'' \uplus \sigma_i) \in \gamma_{\text{pset}}(S_i :: \text{in}(M\sigma'))$, where σ_1 is the instantiation of the received message G as defined in item 3e of Lemma 3.

We now reason on the arbitrary $G' \in \text{int}(\mathcal{C}, G_{\text{lab}})$ of item 3 and we prove that $(S_r :: \text{in}(M\sigma), \sigma_1 \uplus \sigma'') \in \gamma_{\text{pset}}(S_r :: \text{in}_{\text{com}}(G', M\sigma'))$. By definition of γ_{net} , $(S_r, \sigma_r) \in \gamma_{\text{pset}}(S_r)$. By item 3e of Lemma 3, it is easy to verify that the equality constraints introduced by INTER-BIND are preserved by the instantiation σ_1 of the received message. The equality constraints introduced by INTER-MATCH are satisfied as well, since the pattern-matched name n has the same instantiation in both σ_i and σ_1 , the instantiation of the corresponding restriction is the same both in S_r and S_i and so, by definition of γ_{net} , the actions occurring in the same thread and preceding the restriction are instantiated in the same way.

Notice that $S_r \in \text{paths}(\mathcal{C}, \text{out}(G'').Q_r, \emptyset)$ does not necessarily mean that $S_r \in \text{paths}(\mathcal{C}, \text{out}(G'').Q_r, \{\text{in}(M).P \xrightarrow{\text{com}(G)} Q\})$. However, by Lemma 4, we can take some $S' \in \text{paths}(\mathcal{C}, Q, \emptyset)$ such that $S_r \preceq S'$. By Definition 11, S' satisfies the equality constraints in S_r , as desired. By repeating the same reasoning for the other integer components, we get the result.

Condition (ii) This condition is trivially satisfied by considering the substitution $\sigma_1 \uplus \sigma'' \uplus \sigma_i$.

Thus $\langle s :: \text{in}(G), P\sigma \rangle \in \gamma_{\text{net}}(\mathcal{C})$, as desired.

The proof for the replication is straightforward since $\langle s, !P \rangle \rightarrow \langle s, P!P \rangle$ and $\text{threads}(P) = \text{threads}(P!P)$. The proof for the inductive step follows straightforwardly from the induction hypothesis.

Finally, we give the theorems stating that causal graphs constitute an abstract model where secrecy and authenticity can be soundly verified. The next theorem says that if \mathcal{C} is the causal graph associated with P and \mathcal{C} 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} a cycle-invariant causal graph s.t. $\mathcal{C} = \text{graph}(P)$. If \mathcal{C} guarantees the secrecy of v then P guarantees the secrecy of any $v \in \gamma_{\text{frm}}(v)$.

Proof. Straightforwardly by Lemma 3 and Theorem 2.

As mentioned in Section 2.3, we assume that each bound name in the protocol specification has a distinct canonical representative; this helps verifying the secrecy of messages as, fixed a protocol session, every message has a different abstraction in the causal graph. 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. The next theorem states that causal graphs are a safe abstract interpretation of authenticity.

Theorem 4 (Authenticity) Let P be a process and \mathcal{C} a cycle-invariant causal graph s.t. $\mathcal{C} = \text{graph}(P)$.

- 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.

Proof. Directly by Lemma 5 and Theorem 2. Notice that since P is nonce linear, then for every $s \in T(P)$, $\text{end}_n^i(I, J, M) \in s$ and $\text{end}_{n'}^i(I', J', M') \in s$ imply $n \neq n'$, thus getting the desired injectivity property.

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.

D Variable Instantiation with Ciphertexts

In this section we relax the assumption that variables are only instantiated by names. The reason for this assumption is preventing the number of nodes in the causal graphs from diverging, which happens if variables are instantiated with a possibly infinite number of ciphertexts. An effective solution to this problem is to guarantee that the number of ciphertexts generated by trusted principals is finite and to abstract away from the ones generated by the environment. We introduce the new syntactic category of ciphertext variables, ranged over by ξ, ϕ, ψ . These variables can only be instantiated by ciphertexts: as mentioned in the paper, we assume that names and ciphertexts are tagged so as to distinguish them. If ciphertext variables are not encrypted within the process, then it is easy to see that the number of ciphertexts generated by the process is finite. We do not find this over-approximation too inconvenient since nested encryptions are expensive and typically avoided and, in fact, in most protocols principals simply forward ciphertexts since the binding between different ciphertexts is given by some secret such as a nonce or a key (e.g., Kerberos [28]). Furthermore, this approximation does not rule out protocols where a principal generates a nested encryption that does not contain any ciphertext (e.g., Needham Schroeder symmetric-key [27]).

In order to abstract away from the ciphertexts generated by the environment, the special message \mathcal{E} now abstracts over all terms possibly known to the environment, not only the initially known ones.

D.1 ρ -spi calculus and causal graphs

We start by extending function *bind* in order to allow for the instantiation of ρ -spi variables with either names or ciphertexts.

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

Second, we do the same for function $bind_{abs}$:

$$\begin{aligned}
bind_{abs}(v_1, v_2) &= v_2 && \text{if either } [v_1] = [v_2] \\
&&& \text{or } ([v_1] = \mathcal{E} \wedge [v_2] \in \{l, k_1^+, k_{IE}, k_{EI}, k_E^-, \mathcal{E}\}) \\
&&& \text{or } ([v_2] = \mathcal{E} \wedge [v_1] \in \{l, k_1^+, k_{IE}, k_{EI}, k_E^-, \mathcal{E}\}) \\
bind_{abs}(x, a) &= a_{(x)} \\
bind_{abs}(\xi, \{G\}_v) &= \{G\}_v \sigma && \sigma \text{ maps each name } a \text{ in } \{G\}_v \text{ to } a_{(\xi_i)} \text{ for some fresh } \xi_i \\
bind_{abs}((M, M'), (G, G')) &= (G_1, G_2) && \text{if } bind_{abs}(M, G) = G_1 \wedge bind_{abs}(M', G') = G_2 \\
bind_{abs}(\{M\}_v, \{G\}_{v'}) &= \{G_1\}_{v_2} && \text{if } bind_{abs}(M, G) = G_1 \wedge \\
&&& (bind_{abs}(v, v') \neq \uparrow \vee bind_{abs}(v, v') \neq \uparrow) \\
bind_{abs}(M, G) &= \uparrow && \text{otherwise}
\end{aligned}$$

Notice that each name a in the ciphertext replacing variable ξ is mapped to $a_{(\xi_i)}$, for some fresh variable ξ_i . This allows for tracking session messages within the received ciphertext. Function $subst$ is extended accordingly so as to yield a substitution mapping each ciphertext variable to the corresponding ciphertext. For preventing the number of nodes from diverging, we require in function $graph$ that ciphertext variables are only instantiated by ciphertexts whose encryption key is unknown to the environment. This way we prevent variables from being instantiated by the infinite number of ciphertexts generated by the environment, yet allowing for their instantiation by the finite number of ciphertexts generated by trusted processes.

$$threads(P) \subseteq \mathcal{N}$$

$$\begin{aligned}
p.P \in \mathcal{N} \wedge Q \in threads(P) \wedge p \neq in(\cdot) \\
\Rightarrow Q \in \mathcal{N} \wedge p.P \rightarrow Q \in \mathcal{E}
\end{aligned}$$

$$\begin{aligned}
in(M).P \in \mathcal{N} \wedge output(\mathcal{N}) \vdash G \wedge bind_{abs}(M, G) = G' \wedge \sigma = subst(G') \wedge Q \in threads(P\sigma) \wedge \\
G_1 \in int(\mathcal{N}, G') \wedge \{P_1\} = index(G_1) \wedge \{G_2\}_{v_2} \in range(\sigma) \Rightarrow output(\mathcal{N}) \not\vdash v_2 \\
\Rightarrow \{Q, com(G')\} \subseteq \mathcal{N} \wedge \{in(M).P \rightarrow com(G'), P_1 \xrightarrow{G_1} com(G'), com(G') \rightarrow Q\} \subseteq \mathcal{E}
\end{aligned}$$

Finally, we need to refine the concretization functions γ_{trm} and γ_{env} . Their definition is the same as in Table 4, apart from the concretization of \mathcal{E} :

$$\begin{aligned}
\gamma_{trm}(\mathcal{C}, \mathcal{E}) &= ID \cup \mathcal{K}^+ \cup \mathcal{M}_E \cup \mathcal{K}_E \cup \{G \mid G \in \gamma_{env}(\mathcal{C}, G) \wedge output(nodes(\mathcal{C})) \vdash G \wedge G \neq \mathcal{E}\} \\
\gamma_{env}(\mathcal{C}, \mathcal{E}) &= ID \cup \mathcal{K}^+ \cup \mathcal{M}_E \cup \mathcal{K}_E \cup \{G \mid G \in \gamma_{env}(\mathcal{C}, G) \wedge output(nodes(\mathcal{C})) \vdash G \wedge G \neq \mathcal{E}\}
\end{aligned}$$

D.2 Proofs of soundness and safety

Lemma 1 is modified according to the definition of γ_{trm} .

Lemma 6 (Term Concretization). For every term G and causal graph \mathcal{C} , $\gamma_{\text{trm}}(\mathcal{C}, G) \subseteq \gamma_{\text{env}}(\mathcal{C}, G)$.

Proof. The proof is similar to the one of Lemma 1.

In the following, we say that a term M is \mathcal{E} -free if and only if $v \in \text{payload}(M)$ implies $\lfloor v \rfloor \neq \mathcal{E}$, where payload is defined as follows:

$$\text{payload}(M) \triangleq \begin{cases} \{u\} & \text{if } M = u \\ \text{payload}(M_1) \cup \text{payload}(M_2) & \text{if } M = (M_1, M_2) \\ \text{payload}(M') & \text{if } M = \{M'\}_u \end{cases}$$

Furthermore, we say that a term G' is a \mathcal{E} -unfolding of G in \mathcal{C} , written $G' \in \text{unfold}(\mathcal{C}, G)$, if G' is obtained from G by replacing some occurrences of \mathcal{E} in G with some message $\lceil G'' \rceil$ such that $\text{output}(\text{nodes}(\mathcal{C})) \vdash G''$. Lemma 7 is similar to Lemma 2: the only difference is that we require the abstract term M to be \mathcal{E} -free. Notice that input patterns in causal graph generated from a ρ -spi process are always \mathcal{E} -free.

Lemma 7 (Binding). Let \mathcal{C} be a causal graph, G and M ρ -spi terms, G and M abstract terms such that $G \in \gamma_{\text{env}}(\mathcal{C}, G)$, $M \in \gamma_{\text{trm}}(M)$, $\text{bind}(M, G) \neq \uparrow$ and M is \mathcal{E} -free. Then $\exists G' \in \text{unfold}(\mathcal{C}, G)$ such that $\text{bind}_{\text{abs}}(M, G') = G''$, $\sigma = \text{subst}(G'')$, and $\overline{G} \in \gamma_{\text{trm}}(M\sigma)$.

Proof. By induction on the structure of M .

Base Case We have two cases, depending on whether M is a variable or a name:

- Let $M = x$ and $G = v$. By definition of function γ_{trm} , $M = x$, for some x . We have two cases, depending on whether G is a name or a ciphertext. If G is a name, then either $G = v$, for some v such that $n \in \gamma_{\text{env}}(\mathcal{C}, v)$, or $G = \mathcal{E}$. The thesis derives directly from the definition of function bind_{abs} . If G is a ciphertext, then the reasoning is similar.
- Let $M = a$ and $G = a$. We have two cases, depending on whether $G = v$, for some v such that $a \in \gamma_{\text{env}}(\mathcal{C}, v)$, or $G = \mathcal{E}$. In the former case, the proof is as in Lemma 2. Since $n \in \gamma_{\text{env}}(\mathcal{C}, G)$, if $G = \mathcal{E}$, then there exists $v \in \text{unfold}(\mathcal{C}, \mathcal{E})$ such that $n \in \gamma_{\text{trm}}(\mathcal{C}, v)$, as desired.

Inductive Step We show the proof for the case where M is a ciphertext in that the remaining cases are similar.

- Let $M = \{M'\}_{a_1}$ and $G = \{G'\}_{a_2}$, with $\overline{a_1} = a_2$. Since M is \mathcal{E} -free, by definition of γ_{trm} there exist M', v_1, σ'' such that $M = \{M'\}_{v_1}$ and $M \in \gamma_{\text{trm}}(\mathcal{C}, M)$. We have two cases, depending on whether G is a ciphertext or \mathcal{E} . Suppose that there exist G', v_2 such that $G = \{G'\}_{v_2}$, $G' \in \gamma_{\text{env}}(\mathcal{C}, G')$ and $a_2 \in \gamma_{\text{trm}}(\mathcal{C}, v_2)$. (If $G = \mathcal{E}$, then there exists $G' \neq \mathcal{E}$ such that $G \in \gamma_{\text{env}}(\mathcal{C}, G')$ and thus G' is a ciphertext as well.) It is easy to see that if $\overline{a_1} = a_2$ then either $\text{bind}_{\text{abs}}(\overline{v_1}, v_2) \neq \uparrow$ or $\text{bind}_{\text{abs}}(v_1, \overline{v_2}) \neq \uparrow$. By induction hypothesis, there exists $G'' \in \text{unfold}(\mathcal{C}, G')$ such that $\text{bind}_{\text{abs}}(M', G'') = G'''$, $\sigma' = \text{subst}(G''')$ and $G \in \gamma_{\text{trm}}(\mathcal{C}, M'\sigma')$. The thesis follows directly from the definition of bind_{abs} and γ_{trm} .

Since processes do not encrypt ciphertexts received from the network, it is easy to prove that if $\text{output}(\text{nodes}(\mathcal{C})) \vdash G$ then $\text{output}(\text{nodes}(\mathcal{C})) \vdash G'$ for every $G' \in \text{unfold}(\mathcal{C}, G)$. The rest of the proof is similar to the one in Section C.