The Sizes of Skeletons: Decidable Cryptographic Protocol Authentication and Secrecy Goals*

Joshua D. Guttman and F. Javier Thayer

The MITRE Corporation
guttman, jt@mitre.org

Abstract. We show how to *collapse* executions of a cryptographic protocol, when they contain behaviors that we regard as redundant. Moreover, executions containing sufficiently many local runs necessarily contain redundant behaviors, if they have limited numbers of fresh values. Since precise authentication and secrecy assertions are explicit about which values must be assumed to be fresh, it follows that these assertions are decidable.
We formalize these notions within the strand space framework, introducing the notion of a *skeleton*, a collection of behaviors of the regular (non-penetrator) participants. *Homomorphisms* between skeletons express natural relations relevant to protocol analysis.

1 Introduction

It is accepted that the cryptographic protocol problem is undecidable [3]. To find decidable subproblems, one may restrict the behaviors of principals to permit only finitely many runs of the protocol roles, or one may require principals to stop when they have jointly used up a finite budget of fresh random values (nonces). These limitations seem artificial, and unmotivated by protocol behavior. Alternatively, one may consider protocols that never send syntactically similar messages in different situations [1, 10]. Many natural protocols meet this condition, and other protocols can be adapted to it by adding tags that distinguish encrypted units generated at different points in the protocol.

However, a question remains whether the original undecidability result is too pessimistic. Perhaps there exists a class of problems, forming a reasonable set of goals for protocol analysis to resolve, which is in fact decidable for all cryptographic protocols, regardless of the forms of the messages used in those protocols. One motivation for the present paper is to answer this question affirmatively.

The paper also has another motivation. This is to introduce the notion of a *skeleton* (Definition 3), together with homomorphisms between skeletons

* Supported by the MITRE-Sponsored Research program.

(Definition 8). A skeleton provides partial information about the regular (non-penetrator) behavior in some set of possible executions. A homomorphism is an information-preserving map between skeletons. Skeletons and homomorphisms form a category, and much protocol analysis can be regarded as an exploration of properties of this category [2]. In this paper, we use operations on skeletons to show that if a protocol execution involves many runs of the protocol roles, but only a small number of nonces, then there is a smaller execution that is equivalent in a certain sense (Theorem 3). It follows that if there is a counterexample to a formula expressed in a certain fragment of a first order language, then there is also a counterexample using a limited number of runs of the protocol roles. Since there are only finitely many essentially different executions of limited size, the formulas are decidable (Theorem 5). We include proof sketches for most of the propositions below.

Theorem 4 probably also follows from the limited-nonce decidability results. Moreover, Theorem 3 owes something to Heather and Schneider's [6, 7]. However, the results do not appear to have been known previously, and they flow naturally from the skeletons-and-homomorphisms method. We offer them as an introduction to that method, which appears to us more broadly useful [2].

2 Terms, Strands, and Bundles

Terms form a free algebra A, built from atomic terms via constructors. The atomic terms are partitioned into the types *principals*, *texts*, *keys*, and *nonces*. There is an inverse operator defined on keys. Atoms are regarded as indeterminates (variables), and are written in italics (e.g. a, N_a, K^{-1}). We assume A contains infinitely many atoms of each type.

The terms in the algebra A are freely built up from atoms using the operations of *tagged concatenation* and *encryption*. The tags are chosen from a set of constants written in sans serif font (e.g. $\mathsf{tag}, \mathsf{call}$). The tagged concatenation using tag tag of t_0 and t_1 is written $\mathsf{tag} \mathbin{\hat{}} t_0 \mathbin{\hat{}} t_1$; there is a distinguished tag null, and the tagged concatenation using tag null of t_0 and t_1 is written $t_0 \mathbin{\hat{}} t_1$. The encryption operator takes a term t and a key K, and yields a term as result written $\{|t|\}_K$. In the present formulation the second argument to an encryption is always an atomic key.

Fix some choice of algebra A for the remainder of this paper. *Replacements* are defined to have only atoms in their range.

Definition 1 (Replacement, Application). A *replacement* is a function α mapping atoms to atoms, such that (1) for every atom a, $\alpha(a)$ is an atom of the same type as a, and (2) for every key K, $K^{-1} \cdot \alpha = (K \cdot \alpha)^{-1}$.

The *application* of a replacement α to terms t, written $t \cdot \alpha$, is defined to be the homomorphism on terms extending α's action on atoms. More explicitly, if $t = a$ is an atom, then $a \cdot \alpha = \alpha(a)$; and:

$$(\mathsf{tag} \mathbin{\hat{}} t_0 \mathbin{\hat{}} t_1) \cdot \alpha = \mathsf{tag} \mathbin{\hat{}} (t_0 \cdot \alpha) \mathbin{\hat{}} (t_1 \cdot \alpha)$$
$$(\{|t|\}_K) \cdot \alpha = \{|t \cdot \alpha|\}_{K \cdot \alpha}$$

We let replacement application distribute through pairing and sets.Thus, $(x, y) \cdot \alpha = (x \cdot \alpha, y \cdot \alpha)$, and $S \cdot \alpha = \{x \cdot \alpha \colon x \in S\}$. If $x \notin \mathsf{A}$ is a simple value such as an integer or a symbol, then $x \cdot \alpha = x$.

Definition 2 (Strand Spaces). A *direction* is one of the symbols $+, -$. A *directed term* is a pair (d, t) with $t \in \mathsf{A}$ and d a direction, normally written $+t, -t$. $(\pm\mathsf{A})^*$ is the set of finite sequences of directed terms.

A *strand space* over A is a structure containing a set Σ and two mappings: a trace mapping $\mathsf{tr} : \Sigma \to (\pm\mathsf{A})^*$ and a replacement application operator $(s, \alpha) \mapsto s \cdot \alpha$ such that $\mathsf{tr}(s \cdot \alpha) = (\mathsf{tr}(s)) \cdot \alpha$ and $s \cdot \alpha = s' \cdot \alpha$ implies $s = s'$.

Message transmission has positive direction $+$, and reception has negative direction $-$. The conditions ensure that $\cdot$ commutes with tr and that $\cdot$ does not identify distinct strands.

Some additional definitions, including the subterm relation $\sqsubset$ and the penetrator strands (Definition 16), are in Appendix A. Strands that are not penetrator behaviors are called *regular strands*. An important consequence of Definition 16 is that penetrator strands are invariant under replacement (Proposition 17). By a *node* we mean a pair $n = (s, i)$ where $i \leq \mathsf{length}(\mathsf{tr}(s))$; the *direction* and *term* of n are the direction and term of $\mathsf{tr}(s)(i)$ respectively. We prefer to write $s \downarrow i$ for the node $n = (s, i)$. The set $\mathcal{N}$ of all nodes forms a directed graph $\langle \mathcal{N}, (\rightarrow \cup \Rightarrow) \rangle$ together edges $n_1 \rightarrow n_2$ for communication and $n_1 \Rightarrow n_2$ for succession on the same strand (Definition 15). A *bundle* is a subgraph of $\langle \mathcal{N}, (\rightarrow \cup \Rightarrow) \rangle$ for which the edges are causally well-founded, expressing a possible execution (Definition 17).

Proposition 1 (Bundles preserved by replacement). *If $\mathcal{B}$ is a bundle and α is a replacement, then $\mathcal{B} \cdot \alpha$ is a bundle.*

Proof. By Definition 1, $\mathcal{B} \cdot \alpha$ is a graph, and moreover $\mathcal{B} \cdot \alpha$ is isomorphic to $\mathcal{B}$ by the condition (Definition 2) that $s \cdot \alpha = s' \cdot \alpha$ implies $s = s'$. Moreover, if $n \in \mathcal{B}$, then $n \cdot \alpha$ agrees with it in direction and $\mathrm{term}(n \cdot \alpha) = \mathrm{term}(n) \cdot \alpha$. Hence, the bundle conditions (Definition 17) are met in $\mathcal{B} \cdot \alpha$.

We say that t *originates* on n (Appendix A, Definition 15) when t is transmitted at n but was neither received nor transmitted earlier on the same strand. Keys that originate nowhere in a bundle are definitely uncompromised. They may still be used in the bundle, because with our definition of subterm (Definition 15, Clause 1), the encryption key K is not a subterm of $\{|t|\}_K$, unless it was a subterm of t. Values that originate at just one node are *fresh* and suited for use as nonces or (if uncompromised) as session keys.

Proposition 2. *Suppose S is a set of nodes and α is a replacement. (1) If for all a such that $a \cdot \alpha = a_0$, a is non-originating in S, then a_0 is non-originating in $S \cdot \alpha$. (2) If there is no $a' \neq a$ such that $a' \cdot \alpha = a \cdot \alpha$, and a is uniquely originating in S, then $a_0 = a \cdot \alpha$ is uniquely originating in $S \cdot \alpha$.*

Proof. (1) Suppose $a_0 \sqsubset \mathrm{term}(n \cdot \alpha)$ where n is a positive node in S, then $a \sqsubset \mathrm{term}(n)$ and by assumption a is non-originating, so $a \sqsubset \mathrm{term}(n')$ for some $n' \Rightarrow^* n$. Thus $a_0 \sqsubset \mathrm{term}(n' \cdot \alpha)$ and a_0 is non-originating.

(2) Omit the node on which a originates, and apply the first assertion.

3 Preskeletons and Skeletons

A preskeleton is potentially the regular (non-penetrator) part of a bundle or of some portion of a bundle. It is annotated with some additional information, indicating order relations among nodes, uniquely originating atoms, and non-originating atoms. We say that an atom a *occurs* in a set N of nodes if for some $n \in \mathsf{N}$, $a \sqsubset \mathrm{term}(n)$. A key K is *used* in N if for some $n \in \mathsf{N}$, $\{|t|\}_K \sqsubset \mathrm{term}(n)$.

Definition 3. A four-tuple $\mathbb{A} = (\mathsf{N}, \preceq, \mathsf{non}, \mathsf{unique})$ is a *preskeleton* if:

1. N is a finite set of regular nodes; $n_1 \in \mathsf{N}$ and $n_0 \Rightarrow^+ n_1$ implies $n_0 \in \mathsf{N}$;
2. $\preceq$ is a partial ordering on N such that $n_0 \Rightarrow^+ n_1$ implies $n_0 \preceq n_1$;
3. non is a set of keys such that $K \in \mathsf{non}$ implies K does not occur in N, but either K or K^{-1} is used in N;
4. unique is a set of atoms such that $a \in \mathsf{unique}$ implies a occurs in N.

A preskeleton $\mathbb{A}$ is a *skeleton* if in addition:

$4'$. $a \in \mathsf{unique}$ implies a originates at no more than one node in N.

We select components of a preskeleton using subscripts. For instance, if $\mathbb{A} = (\mathsf{N}, R, S, S')$, then $\preceq_{\mathbb{A}}$ means R and $\mathsf{unique}_{\mathbb{A}}$ means S'. We write $n \in \mathbb{A}$ to mean $n \in \mathsf{N}_{\mathbb{A}}$, and we say that a strand s is in $\mathbb{A}$ when at least one node of s is in $\mathbb{A}$. The $\mathbb{A}$-height of s is the number of nodes of s in $\mathbb{A}$. By Clauses 3 and 4, $\mathsf{unique}_{\mathbb{A}} \cap \mathsf{non}_{\mathbb{A}} = \emptyset$. Bundles correspond to certain skeletons:

Definition 4. Bundle $\mathcal{B}$ *realizes* skeleton $\mathbb{A}$ if (1) the nodes of $\mathbb{A}$ are precisely the regular nodes of $\mathcal{B}$; (2) $n \preceq_{\mathbb{A}} n'$ just in case $n, n' \in \mathsf{N}_{\mathbb{A}}$ and $n \preceq_{\mathcal{B}} n'$; (3) $K \in \mathsf{non}_{\mathbb{A}}$ just in case K or K^{-1} is used in $\mathsf{N}_{\mathbb{A}}$ but K occurs nowhere in $\mathcal{B}$; (4) $a \in \mathsf{unique}_{\mathbb{A}}$ just in case a originates uniquely in $\mathcal{B}$.

The *skeleton* of $\mathcal{B}$, written $\mathsf{skeleton}(\mathcal{B})$, is the skeleton that it realizes.

If $\mathcal{B}$ is a bundle, then there is a unique skeleton that it realizes. By condition (4), $\mathcal{B}$ does not realize $\mathbb{A}$ if $\mathbb{A}$ is a preskeleton but not a skeleton.

We also want to view realizability more locally. A negative node in a preskeleton is realizable when the adversary can derive the message received, using terms transmitted on earlier positive nodes.

Definition 5 (Penetrator web). Let $G = \langle \mathcal{N}_G, (\rightarrow_G \cup \Rightarrow_G) \rangle$ be a finite acyclic subgraph of $\langle \mathcal{N}, (\rightarrow \cup \Rightarrow) \rangle$ such that $\mathcal{N}_G$ consists entirely of penetrator nodes. G is a *penetrator web* with support S and result R if S and R are sets of terms and moreover:

1. If $n_2 \in \mathcal{N}_G$ is negative, then either $\mathrm{term}(n_2) \in S$ or there is a unique n_1 such that $n_1 \rightarrow_G n_2$.
2. If $n_2 \in \mathcal{N}_G$ and $n_1 \Rightarrow n_2$ then $n_1 \Rightarrow_G n_2$.
3. For each $t \in R$, either $t \in S$ or there is some positive $n \in \mathcal{N}_G$ such that $\mathrm{term}(n) = t$.

Proposition 3. *If $\mathcal{B}$ is a bundle and $n \in \mathcal{B}$ is negative, let G be the set of penetrator nodes m of $\mathcal{B}$ such that $m \preceq_{\mathbb{A}} n$ and S the set of terms $term(m)$ for m regular and positive and $m \preceq_{\mathbb{A}} n$. Then G is a penetrator web with support S and result $\{term(n)\}$.*

Definition 6 (Realizable node). If $\mathbb{A}$ is a preskeleton and $n \in \mathbb{A}$ is negative, then a penetrator web G with support S and result R *realizes node* n in $\mathbb{A}$ if

1. $\mathrm{term}(n) \in R$;
2. $S \subset \{\mathrm{term}(m) \colon m \preceq_{\mathbb{A}} n \text{ and } m \text{ is positive}\}$;
3. a originates in G implies $a \notin \mathsf{non}_{\mathbb{A}}$;
4. a originates in G and a originates in $\mathbb{A}$ implies $a \notin \mathsf{unique}_{\mathbb{A}}$.

Node n is *realizable in* $\mathbb{A}$ if there is a penetrator web G that realizes it.

Proposition 4. *If $\mathcal{B}$ is a bundle with $n \in \mathcal{B}$ negative, then there is a subgraph G_n of $\mathcal{B}$ such that G_n realizes n in $\mathsf{skeleton}(\mathcal{B})$.*

Skeleton $\mathbb{A}$ is realizable if and only if every negative $n \in \mathbb{A}$ is realizable in $\mathbb{A}$.

Proof. The first assertion follows from Proposition 3. The second assertion holds (left-to-right) by the previous assertion. Right-to-left, it follows by taking the union of the penetrator webs, identifying any minimal penetrator M, K-nodes originating the same term.

This last assertion holds only for *skeletons* $\mathbb{A}$, because a non-skeleton is never realizable. Inspired by Proposition 2, we define:

Definition 7. A replacement α *respects origination in* $\mathbb{A}$ just in case (1) for all a, a', if $a \in \mathsf{non}_{\mathbb{A}}$ and $a \cdot \alpha = a' \cdot \alpha$ then $a' \in \mathsf{non}_{\mathbb{A}}$; and (2) for all a, a', if $a \in \mathsf{unique}_{\mathbb{A}}$ and $a \cdot \alpha = a' \cdot \alpha$, then $a = a'$.

If α is injective, then it respects origination. By Proposition 17, being a penetrator web is invariant under replacement. Using Definitions 5 and 7, we have:

Proposition 5. *If n is realizable in preskeleton $\mathbb{A}$ and α respects origination in $\mathbb{A}$, then $n \cdot \alpha$ is realizable in $\mathbb{A} \cdot \alpha$. If skeleton $\mathbb{A}$ is realizable and α respects origination, then $\mathbb{A} \cdot \alpha$ is realizable.*

Proposition 6. *It is decidable whether node n is realizable in skeleton $\mathbb{A}$. Hence, it is decidable whether $\mathbb{A}$ is realizable.*

Proof. If $\mathbb{A}$ is realizable, then there is a *normal* bundle $\mathcal{B}$ that realizes it [5]. Thus, when $n \in \mathbb{A}$ is realizable, we will use only subterms of $\{\mathrm{term}(m) \colon m \preceq_{\mathbb{A}} n \text{ and } m \text{ positive}\} \cup \{\mathrm{term}(n)\}$, of which there are only finitely many.

Proposition 6 is well-known, e.g. [8], with a stronger form, where replacements may carry variables to terms, not just atoms. Whether $\mathbb{A}$ may be embedded in a realizable skeleton $\mathbb{A}'$ is a different matter; it is undecidable [3] for definitions such as our Definition 9.

4 Collapsing Skeletons

We show next how to collapse preskeletons without destroying realizability. When α unifies two strands s, s' in $\mathbb{A}$, and α respects origination in $\mathbb{A}$, then we can equate $s \cdot \alpha$ with $s' \cdot \alpha$ in $\mathbb{A} \cdot \alpha$. Identifying them, while leaving other strands distinct, yields a realizable preskeleton $\mathbb{A}'$, if $\mathbb{A}$ was realizable.

These operations on preskeletons use homomorphisms in the following sense.

Definition 8. Let $\mathbb{A}_0, \mathbb{A}_1$ be preskeletons, α a replacement, $\phi\colon \mathsf{N}_{\mathbb{A}_0} \to \mathsf{N}_{\mathbb{A}_1}$. $H = [\phi, \alpha]$ is a *homomorphism* if

1. $\mathrm{term}(\phi(n)) = \mathrm{term}(n) \cdot \alpha$ for all $n \in \mathbb{A}_0$
1′. $m \Rightarrow \phi(n')$ iff $m = \phi(n)$ where $n \Rightarrow n'$
2. $n \preceq_{\mathbb{A}_0} m$ implies $\phi(n) \preceq_{\mathbb{A}_1} \phi(m)$
3. $\mathsf{non}_{\mathbb{A}_0} \cdot \alpha \subset \mathsf{non}_{\mathbb{A}_1}$
4. $\mathsf{unique}_{\mathbb{A}_0} \cdot \alpha \subset \mathsf{unique}_{\mathbb{A}_1}$

We write $H\colon \mathbb{A}_0 \mapsto \mathbb{A}_1$ to mean that H is a homomorphism from $\mathbb{A}$ to $\mathbb{A}'$.

When $a \cdot \alpha = a \cdot \alpha'$ for every a used or uttered in $\mathsf{dom}(\phi)$, then $[\phi, \alpha] = [\phi, \alpha']$; i.e., $[\phi, \alpha]$ is the equivalence class of pairs under this relation.

When a homomorphism identifies nodes, only one of them needs to be realizable:

Proposition 7. *Let $\mathbb{A}, \mathbb{A}'$ be preskeletons, and let $H = [\phi, \alpha]\colon \mathbb{A} \mapsto \mathbb{A}'$, where α respects origination in $\mathbb{A}$. (1) If there exists any $n \in \mathbb{A}$ such that n is realizable in $\mathbb{A}$ and $\phi(n) = m$, then m is realizable in $\mathbb{A}'$.*

(2) Suppose $\mathbb{A}'$ is a skeleton. If for every $m \in \mathbb{A}'$ there exists an $n \in \mathbb{A}$ such that n is realizable in $\mathbb{A}$ and $\phi(n) = m$, then $\mathbb{A}'$ is realizable.

Proof. (1) holds by Proposition 5. (2) holds by (1) and Proposition 4.

We recall that any partial order $\leq$ (or indeed any reflexive, transitive relation) can be regarded as a graph G, in which there is an edge $x \rightharpoonup y$ just in case $x \leq y$ and for all z, $x \leq z \leq y$ implies $x = z$ or $z = y$. When we say that a node n *immediately precedes* m in $\mathbb{A}$, we mean that $n \rightharpoonup m$ in the graph $G(\mathbb{A})$ generated from $\preceq_{\mathbb{A}}$. In $G(\mathbb{A})$, there may not be an arrow $n_0 \rightharpoonup n_1$ when $n_0 \Rightarrow n_1$; this happens in case n_0 is positive, n_1 is negative, and there is at least one node m not on this strand such that $n_0 \preceq_{\mathbb{A}} m \preceq_{\mathbb{A}} n_1$. If $\mathbb{A} = \mathsf{skeleton}(\mathcal{B})$, this is the only case in which $n_0 \Rightarrow n_1$ but there is no arrow $n_0 \rightharpoonup n_1$ in G.

An edge $n_0 \rightharpoonup n_1$ in $G(\mathbb{A})$ is *removable* when $n_0 \not\Rightarrow n_1$; homomorphisms cannot change the strand structure between nodes, but they can enrich the order to add back any removable edge. We call a homomorphism $H = [\mathsf{id}, \mathsf{id}]\colon \mathbb{A} \mapsto \mathbb{A}'$ an *order enrichment* when $\mathsf{N}_{\mathbb{A}} = \mathsf{N}_{\mathbb{A}'}$, $\mathsf{non}_{\mathbb{A}} = \mathsf{non}_{\mathbb{A}'}$, and $\mathsf{unique}_{\mathbb{A}} = \mathsf{unique}_{\mathbb{A}'}$. Hence, the only possible difference is that $\preceq_{\mathbb{A}'}$ may extend $\preceq_{\mathbb{A}}$.

Proposition 8. *Suppose that $\mathbb{A}'$ is a preskeleton and S is a set of removable edges in $G(\mathbb{A}')$. There is a preskeleton $\mathbb{A}$ and an order enrichment $H\colon \mathbb{A} \mapsto \mathbb{A}'$ such that $G(\mathbb{A}') \setminus S = G(\mathbb{A})$.*

Theorem 1. *Suppose that s_0, s_1 have heights h_0, h_1 (resp.) in the preskeleton $\mathbb{A}'$, with $h_0 \leq h_1$, and suppose that for all $j \leq h_0$, $term(s_0 \downarrow j) = term(s_1 \downarrow j)$ with the same direction. There exist $\mathbb{A}, \mathbb{A}''$, an order enrichment $H\colon \mathbb{A} \mapsto \mathbb{A}'$, and a homomorphism $H'' = [\phi, \mathsf{id}]\colon \mathbb{A} \mapsto \mathbb{A}''$ such that:*

1. *There is a set S containing only removable edges $n \rightharpoonup m$ for which m lies on s_0 or s_1 and $G(\mathbb{A}') \setminus S = G(\mathbb{A})$;*
2. *ϕ is surjective; $\mathsf{non}_{\mathbb{A}''} = \mathsf{non}_{\mathbb{A}'}$; $\mathsf{unique}_{\mathbb{A}''} = \mathsf{unique}_{\mathbb{A}'}$;*
3. *$\phi(n) = n$ unless $n = s_0 \downarrow j$, for some j with $1 \leq j \leq h_0$;*
4. *$\phi(s_0 \downarrow j) = s_1 \downarrow j$, for all j with $1 \leq j \leq h_0$.*

If $\mathbb{A}''$ satisfies (1)–(4) and $n \in \mathbb{A}'$ is realizable in $\mathbb{A}'$, then $\phi(n)$ is realizable in $\mathbb{A}''$.

Proof. Consider any path p through $G(\mathbb{A}')$ leading from $s_0 \downarrow j$ to $s_1 \downarrow j$ (as in Figure 1), or vice versa; let s be the strand at which p ends. There is an edge $n \rightharpoonup s \downarrow j'$ such that $j' \leq j$, n does not lie on s, and no node of s precedes n along p. Let S be the set of all such edges. Hence $G(\mathbb{A}') \setminus S$ remains acyclic, even when each $s_0 \downarrow j$ is identified with $s_1 \downarrow j$.

Let $\mathsf{N}_{\mathbb{A}''}$ be the subset of $\mathsf{N}_{\mathbb{A}}$ of $\mathbb{A}$ not lying on s_0. Let $\mathsf{non}_{\mathbb{A}''} = \mathsf{non}_{\mathbb{A}'}$ and $\mathsf{unique}_{\mathbb{A}''} = \mathsf{unique}_{\mathbb{A}'}$; let $\preceq_{\mathbb{A}''}$ be $\preceq_{\mathbb{A}}$ restricted to $\mathsf{N}_{\mathbb{A}''}$. Define ϕ by (2,3).

When $i = 1, 2$ and $s_i \downarrow j$ is negative, there is a penetrator web $G_{i,j}$ with result $\mathrm{term}(s_i \downarrow j)$, since $\mathbb{A}'$ is realizable. The support of the $G_{i,j}$ contains only terms on earlier positive nodes n. By acyclicity of $\mathbb{A}'$, for a given j, at most one of $s_i \downarrow j$ and $s_{i'} \downarrow j$ precedes the other; say (e.g.) $s_1 \downarrow j$ does not precede $s_0 \downarrow j$.

Hence, $G_{0,j}$'s support does not contain nodes n such that $s_1 \downarrow j \preceq_{\mathbb{A}'} n$. So, web $G_{0,j}$ is still supported in $\mathbb{A}$, and thus $\mathrm{term}(s_0 \downarrow j)$ is realizable in $\mathbb{A}$. Since the replacement id respects origination, Proposition 7 shows (4).

Corollary 2. *Let $\mathbb{A}'$ be a preskeleton containing nodes of the strands s_0, s_1 up to heights h_0, h_1 (resp.) with $h_0 \leq h_1$. Let α respect origination for $\mathbb{A}'$ and unify s_0, s_1 up to h_0, i.e.*

$$term(s_0 \downarrow j) \cdot \alpha = term(s_1 \downarrow j) \cdot \alpha$$

with the same direction, for each j with $1 \leq j \leq h_0$. Then $\mathbb{A}'$ is an order enrichment of some $\mathbb{A}$ such that $H'' = [\phi, \alpha]\colon \mathbb{A} \mapsto \mathbb{A}''$ where

1. *There is a set S containing only removable edges $n \rightharpoonup m$ for which m lies on s_0 or s_1 and $G(\mathbb{A}') \setminus S = G(\mathbb{A})$;*
2. *ϕ is surjective; $\mathsf{non}_{\mathbb{A}''} = \mathsf{non}_{\mathbb{A}'} \cdot \alpha$; $\mathsf{unique}_{\mathbb{A}''} = \mathsf{unique}_{\mathbb{A}'} \cdot \alpha$;*
3. *ϕ is injective for nodes not lying on s_0, s_1;*
4. *$\phi(s_0 \downarrow j) = \phi(s_1 \downarrow j)$ for all $j \leq h_0$.*

If $\mathbb{A}''$ satisfies (1)–(4) and $n \in \mathbb{A}'$ is realizable in $\mathbb{A}'$, then $\phi(n)$ is realizable in $\mathbb{A}''$.

Proof. Apply first Proposition 5 and then Theorem 1.

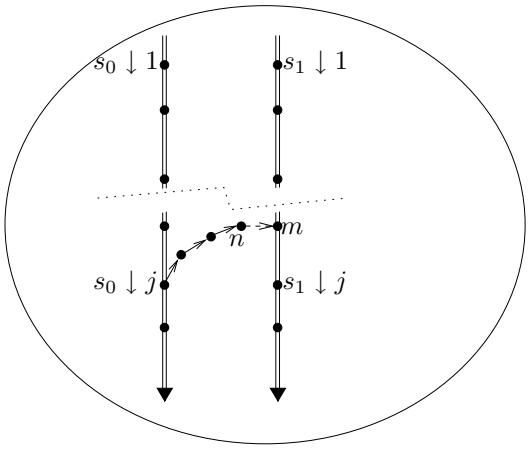

Fig. 1. Removing edge $n \rightharpoonup m$

5 Protocols

Here we introduce the notion of protocol. We prove that if a protocol has a large realizable skeleton using only a small number of non-originating and uniquely originating values, then the skeleton can be collapsed into a smaller skeleton (Theorem 3). For a class of security goals, the smaller skeleton satisfies the same goals as the larger skeleton; we show this in Propositions 12, 16. Since there are only finitely many essentially different skeletons of limited size (Proposition 15), a class of formulas of this language are decidable (Theorem 5).

Definition 9 (Protocol). A *protocol* Π consists of (1) a finite set of strands called the *roles* of Π; (2) for each role $r \in \Pi$, two sets of atoms n_r, u_r giving *origination data* for r; and (3) a number of *key function symbols*, and for each role r a set of 0 or more *key constraints*, i.e. equations involving these function symbols and atoms occurring in r. The *regular strands* of Π, written Σ_Π, are all strands s with $\mathsf{tr}(s) = \mathsf{tr}(r \cdot \alpha)$ for some role $r \in \Pi$.

The origination data n_r, u_r gives values mentioned in r that should be assumed to be non-originating or uniquely originating (respectively) whenever role r is executed. The key constraints give a way to ensure that different parameters are compatible across different strands in the same skeleton or bundle. For instance, one wants to assume that if the same principal executes the same role twice, then it is using the same private decryption key in both runs. In particular, if the principal uses a non-originating private key in one run, then its private key in other runs is also non-originating. The key functions may be assumed injective; an attack that relies on two different principals having chosen the same private key (for instance) has negligible probability of success.

Definition 10 (Skeleton, bundle of a protocol). The *key constraints* of $\mathbb{A}$ are the formulas $\phi \cdot \alpha$ such that ϕ is a key constraint for some role r with $r \cdot \alpha$ in $\mathbb{A}$. $\mathbb{A}$ is a *skeleton for protocol* Π if

1. the strands of $\mathbb{A}$ belong to Σ_Π;
2. if $a \in n_r$ and $a \cdot \alpha$ is used on a node of $r \cdot \alpha$ in $\mathbb{A}$, then $a \cdot \alpha \in \mathsf{non}_{\mathbb{A}}$;
2′. if $a \in u_r$ and $a \cdot \alpha$ occurs in a node of $r \cdot \alpha$ in $\mathbb{A}$, then $a \cdot \alpha \in \mathsf{unique}_{\mathbb{A}}$; and
3. There is a interpretation of the key function symbols by injective functions satisfying all of the key constraints of $\mathbb{A}$.

A bundle $\mathcal{B}$ is a *bundle for* Π if $\mathsf{skeleton}(\mathcal{B})$ is a skeleton for Π.

We will assume that each protocol Π has two *listener roles*. A listener role is a regular strand with a single negative node, where the term received is a single atom. It records the fact that this atom is available on its own, unprotected by encryption, disclosed to the penetrator. HearNonce$[N]$ records the disclosure of nonce N via trace $\langle -N \rangle$; HearKey$[K]$ records the disclosure of key K via trace $\langle -K \rangle$. We refer to a *reduction* when the situation in Corollary 2 holds:

Definition 11 (Reduction). $\mathbb{A}'$ *reduces to* $\mathbb{A}''$ *via* α *in one step*, which we write $\mathcal{R}(\mathbb{A}', \mathbb{A}'', \alpha)$, when there exist $\mathbb{A}, s_0, s_1, h_0, h_1, \phi$ such that

1. s_0, s_1 have $\mathbb{A}'$-heights h_0, h_1 (resp.) with $h_0 \leq h_1$, α respects origination for $\mathbb{A}'$ and unifies terms on s_0, s_1 up to h_0 (with matching direction);
2. $\mathbb{A}'$ is an order enrichment of $\mathbb{A}$ and $H'' = [\phi, \alpha] \colon \mathbb{A} \mapsto \mathbb{A}''$;
3. ϕ is surjective; $\mathsf{non}_{\mathbb{A}''} = \mathsf{non}_{\mathbb{A}'} \cdot \alpha$; $\mathsf{unique}_{\mathbb{A}''} = \mathsf{unique}_{\mathbb{A}'} \cdot \alpha$;
4. ϕ is injective for nodes not lying on s_0, s_1; and
5. $\phi(s_0 \downarrow j) = \phi(s_1 \downarrow j)$ for all $j \leq h_0$.

$\mathbb{A}_0$ *reduces to* $\mathbb{A}_k$ via α when $k = 1$ and $\mathbb{A}_0$ and $\mathbb{A}_1$ are isomorphic, or $\alpha = \alpha_0 \circ \cdots \circ \alpha_{k-1}$, and $\mathcal{R}(\mathbb{A}_i, \mathbb{A}_{i+1}, \alpha_i)$ for each $0 \leq i < k$. We write this $\mathcal{R}^*(\mathbb{A}', \mathbb{A}'', \alpha)$.

Proposition 9. *If* $\mathbb{A}_0$ *reduces to* $\mathbb{A}_\ell$ *and* $\mathbb{A}_0$ *is realizable, then* $\mathbb{A}_\ell$ *is realizable.*

Theorem 3. *Let* Π *be a protocol with* i *roles, each of which has at most* j *parameters. Let* $\mathbb{A}_0$ *be a realizable skeleton for* Π *in which the number of non-originating and uniquely originating values is* k*, i.e.* $|\mathsf{non}_{\mathbb{A}_0} \cup \mathsf{unique}_{\mathbb{A}_0}| = k$*. Then* $\mathcal{R}^*(\mathbb{A}_0 \mathbb{A}_\ell, \alpha)$ *for some realizable* $\mathbb{A}_\ell$ *with at most* $i(j^{k+1})$ *strands.*

Proof. Since there are at most k atoms in $\mathsf{non}_{\mathbb{A}'} \cup \mathsf{unique}_{\mathbb{A}'}$ of any one type, there are at most $k + 1$ values of this type that cannot be unified by a replacement that respects origination for $\mathbb{A}'$. Since each role $r \in \Pi$ has at most j parameters, there are at most j^{k+1} strands of role r that cannot be unified by a replacement respecting origination for $\mathbb{A}'$. As Π contains i roles, there are at most $i(j^{k+1})$ strands such that no two can be unified by a replacement respecting origination for $\mathbb{A}'$. Thus, as long as $\mathbb{A}'$ contains more than this number of strands, we may apply Corollary 2 to obtain a smaller one, preserving realizability.

Applying the ideas of [3], it is undecidable, given a protocol Π, whether a particular parameter of a role $r \in \Pi$ remains secret. That is, given Π, $r \in \Pi$, and a, is there a bundle $\mathcal{B}$ for Π containing a strand $s = r \cdot \alpha$ up to its full height, and a node $n \in \mathcal{B}$ such that $\mathrm{term}(n) = a \cdot \alpha$? Theorem 3 tells us that in the hard choices of Π, the number of values in $|\mathsf{non} \cup \mathsf{unique}|$ increases beyond any k.

6 Security Goals

By Theorem 3, a workable strategy for determining whether a protocol has realizable skeletons of a particular kind is to ensure that unique and non grow far more slowly than the number of strands. In particular, this will be true if for all roles $r \in \Pi$, $n_r = u_r = \emptyset$. We say that Π *does not impose origination assumptions* if this is true.

Is there is any value to protocols that impose no origination assumptions? No interesting conclusions follow if there are no assumptions whatever about origination, but these assumptions need not be imposed by the protocol itself. We may instead define the protocol with $n_r = u_r = \emptyset$, while considering only bundles in which certain values are uniquely originating or non-originating. For instance, in the case of the Needham-Schroeder-Lowe protocol, one can prove [5], letting Π be a representation with $n_r = u_r = \emptyset$:

> Let $\mathcal{B}$ be a bundle for Π with $s \in \text{NSResp}[A, B, N_a, N_b]$ of $\mathcal{B}$-height 3.
> Assume K_A^{-1} is non-originating; N_b is uniquely originating; and $N_a \neq N_b$.
> Then $\mathcal{B}$ contains an $s' \in \text{NSInit}[A, B, N_a, N_b]$ with $\mathcal{B}$-height 3.

The public keys K_A, K_B do not appear here as independent parameters, because they are given by a key function of the principal. By contrasting, the result from the initiator's point of view makes a slightly different assumption:

> Let $\mathcal{B}$ be a bundle for Π with $s \in \text{NSInit}[A, B, N_a, N_b]$ of $\mathcal{B}$-height ≥ 2.
> Assume K_A^{-1}, K_B^{-1} is non-originating and N_a is uniquely originating.
> Then $\mathcal{B}$ contains an $s' \in \text{NSResp}[A, B, N_a, N_b]$ with $\mathcal{B}$-height ≥ 2.

Here it is necessary to assume both private keys K_A^{-1}, K_B^{-1} are uncompromised.

Results of this form are more informative than results where $n_r \neq \emptyset$ or $u_r \neq \emptyset$: one learns that a protocol correctness goal depends only on specific keys or nonces. Neither authentication result depends on the freshness of the other party's key. If instead we were to set $u_r = \{N_b\}$ for the responder role and $u_i = \{N_a\}$ for the initiator role, then this distinctions would be lost: then any bundle containing s, s' would actually have both N_a and N_b uniquely originating.

Putting parameters in n_r, n_i for non-origination is trickier. If the initiator's role has $n_i = \{K_A^{-1}, K_B^{-1}\}$, then a responder C is trusting the initiator A to connect only with regular parties B. This assumption, suggested in [9], makes the protocol valid [4, Section 3.12], but is unreasonable in many environments. Hence, explicit style security goals are more flexible, more informative, and in fact more decidable than properties of protocols imposing origination assumptions.

Both authentication and secrecy goals are universally quantified implications, with premises conjoined from formulas of the form:

roles a strand of a particular role and given parameters has $\mathcal{B}$-height j;
origination a parameter is uniquely originating or non-originating in $\mathcal{B}$;
inequality two parameters are not equal.

The conclusions take forms including:

authentication some strand of a given role with parameters shared with the premises has $\mathcal{B}$-height k; other parameters may be existentially bound.

nondisclosure no regular strand sharing parameters with the premises has $\mathcal{B}$-height k; the remaining parameters may be universally bound.

A secrecy goal is a non-disclosure goal in which the regular strand is a listener strand (see Section 5) and $k = 1$. Other nondisclosure goals might state (e.g.) that there is no strand with parameter p (indicating a regular principal) and nonce n; this would express that the protocol cannot disclose n to p. We will show that these formulas are decidable (Corollary 5).

We formalize these security goals in a first order language. The structures (interpretations) for this language are realizable skeletons. For each role $r \in \Pi$, let the distinct atoms occurring in r be $\boldsymbol{a}^r$, i.e. the atoms $a^r_1, \ldots, a^r_k$ for some k. A role r has length ℓ, written $\mathsf{length}(r) = \ell$, if $\mathsf{tr}(r)$ is a sequence of length ℓ. In the language $\mathcal{L}_\Pi$ that we will define, variables range over atoms, and an interpretation is specified by giving a realizable skeleton. There are predicates saying that an atom is uniquely or non-originating, and predicates saying that the skeleton has a strand of role r with height at least m, for each $r \in \Pi$ and m less than its length.

Definition 12. The language $\mathcal{L}_\Pi$ for protocol Π is the first order language with equality and the propositional constants truth and $\mathsf{falsehood}$, the predicates $\mathsf{nonp}(x)$ and $\mathsf{uniquep}(x)$, and for each r, m where r is a role in Π and $m \leq \mathsf{length}(r)$:

$\phi^r_m(x_1, \ldots, x_k)$ a predicate with k arguments if $\boldsymbol{a}^r$ contains k atoms.

The set of variables occurring free in a formula ψ is $\mathsf{fv}(\psi)$.

Thus, the number of non-logical predicates contained in $\mathcal{L}_\Pi$ is $2+\sum_{r\in\Pi} \mathsf{length}(r)$. If $\sigma(y) = \sigma'(y)$ whenever y is a variable other than x, σ is an *x-variant* of σ':

Definition 13. An *$\mathcal{L}_\Pi$-skeleton structure* (or simply a *structure*) $\mathcal{M} = (\mathbb{A}, \sigma)$ is a realizable skeleton $\mathbb{A}$ and an assignment mapping variables x to atoms a.

$\mathcal{M}$ *satisfies* $\phi^r_m(x_1, \ldots, x_k)$ if $\mathbb{A}$ contains a strand s of $\mathbb{A}$-height m such that $\mathsf{tr}(s) = \mathsf{tr}(r{\cdot}\alpha)$, where $a^r_i{\cdot}\alpha = \sigma(x_i)$ for each i. $\mathcal{M}$ satisfies $\mathsf{nonp}(x)$ if $\sigma(x) \in \mathsf{non}_\mathbb{A}$, and it satisfies $\mathsf{uniquep}(x)$ if $\sigma(x) \in \mathsf{unique}_\mathbb{A}$.

Satisfaction for compound formulas is classical; for instance, $(\mathbb{A}, \sigma) \models \forall x\ .\ \psi$ if for every x-variant σ' of σ, $(\mathbb{A}, \sigma') \models \psi$.

7 Formulas Preserved under Reductions

Proposition 10. *Let ψ be a formula of $\mathcal{L}_\Pi$, and let $\mathcal{M} = (\mathbb{A}, \sigma)$ be a structure. It is decidable whether $\mathcal{M}$ satisfies ψ.*

Proof. By induction on the structure of ψ. In the case of the quantifiers, observe that if the set of atoms mentioned (whether occurring as subterms or used as keys) in $\mathbb{A}$ is S, and there are k variables occurring free in ψ, there are essentially at most $|S| + k + 1$ relevantly different choices for a variable x bound by the outermost quantifier. It takes a value in S, or equals one of the k free variables, or neither.

Two structures are *elementary equivalent* for a (first order) language if they satisfy the same formulas of the language. For "reduction," see Definition 11.

Proposition 11. *Suppose that*

$$\mathbb{A} = (\mathsf{N}, \preceq, \mathsf{non}, \mathsf{unique}) \text{ and } \mathbb{A}' = (\mathsf{N}', \preceq', \mathsf{non}, \mathsf{unique})$$

have the same non *and* unique*, and* $\mathbb{A}'$ *reduces to* $\mathbb{A}$ *in one step via the identity replacement* id*. For every* σ*,* $(\mathbb{A}, \sigma)$ *and* $(\mathbb{A}', \sigma)$ *are elementary equivalent for* $\mathcal{L}_{\Pi}$*.*

Proof. For every σ, $(\mathbb{A}, \sigma)$ and $(\mathbb{A}', \sigma)$ satisfy the same open atomic formulas. The property is preserved under propositional connectives, and because it holds for all σ, it is preserved under quantification.

Revise me, please. What if α is not the identity, because it maps some atoms to the same result? Do we know that a reduction via such an α preserves anything? In this section we will consider what follows for formulas that are preserved by reasonable reductions. In the next section, we will identify a class of formulas, sufficient to express the kinds of security goals described in Section 6, that are in fact preserved by these reductions. These are certain closed formulas of the form

$$\forall x_1, \ldots, x_k \,.\, H \supset C$$

where H is a quantifier-free formula that may use all the predicates of $\mathcal{L}_{Pi}$. $C \in \mathcal{L}_{\Pi}$ by contrast may contain quantifiers, but the quantified body is an atomic formula, and specifically a role predicate $\phi^r_m(x_1, \ldots, x_k)$.

We are only interested in reductions that are "well-behaved" on $\mathsf{non}_{\mathbb{A}}$, $\mathsf{unique}_{\mathbb{A}}$, and $\sigma(\mathsf{fv}(C))$, the image, under σ, of the set of free variables of the formula C to be preserved.

Definition 14 (Well-behaved reductions). A reduction from $\mathbb{A}$ to $\mathbb{A}'$ via α is *well-behaved* for a set of atoms X, written $\mathcal{W}^*_X(\mathbb{A}, \mathbb{A}', \alpha)$, iff $\mathcal{R}^*(\mathbb{A}, \mathbb{A}', \alpha)$ and (1) α is a bijection between $\mathsf{non}_{\mathbb{A}}$ and $\mathsf{non}_{\mathbb{A}'}$; (2) α is a bijection between $\mathsf{unique}_{\mathbb{A}}$ and $\mathsf{unique}_{\mathbb{A}'}$; and (3) for all $a \in X$, $a \cdot \alpha = b \cdot \alpha$ implies $b = a$.

$C \in \mathcal{L}^0_{\Pi}$ is *preserved under well-behaved reductions* if $\mathcal{W}^*_X(\mathbb{A}, \mathbb{A}', \alpha)$ implies $(\mathbb{A}, \sigma) \models C$ if and only if $(\mathbb{A}', \sigma \circ \alpha) \models C$.

Proposition 12. *If* $H \in \mathcal{L}_{\Pi}$ *is quantifier-free, then* H *is preserved under well-behaved reductions.*

Proof. For convenience, suppose that H is written using $\vee$ and $\wedge$, where de Morgan's laws have been used to push negations inwards; i.e., whenever $\neg\psi$ is a subformula of H, then ψ is an atomic formula. We argue by induction on the structure of H. Atomic formulas are preserved all reductions; negated atoms $\mathsf{nonp}(x)$ and respectively $\mathsf{uniquep}(x)$ are preserved by clauses (1,3) and respectively (2,3) of the definition; and negated role predicates and negated equalities are preserved by clause (3). The induction through $\vee$ and $\wedge$ is immediate.

When V is a set of variables, let us write $\sigma \sim_{\mathsf{V}} \sigma'$ to mean that σ and σ' differ on V by a bijective replacement; i.e., there is a replacement α that $\alpha(\sigma(v)) = \sigma'(v)$ for all $v \in \mathsf{V}$, and α is a bijection on all atoms.

Proposition 13. *If V is finite, then there is an integer M such that: any set S of variable assignments in which $\sigma \sim_{\mathsf{V}} \sigma'$ implies $\sigma = \sigma'$ for all $\sigma, \sigma' \in \mathsf{V}$ has cardinality $|S| \leq M$.*

Hence, for finite V, there is a finite maximal set $\mathcal{S}_{\mathsf{V}}$ such that $\sigma, \sigma' \in \mathcal{S}_{\mathsf{V}}$ and $\sigma \sim_{\mathsf{V}} \sigma'$ implies $\sigma = \sigma'$.

Proof. If σ and σ' agree on the type of atom assigned to each variable in V, and they agree on when two variables have the same atom assigned, then $\sigma \sim_{\mathsf{V}} \sigma'$. There are only finitely many ways to avoid these kinds of agreement for finite V.

Proposition 14. *Let ψ be a formula of $\mathcal{L}_{\Pi}$ with $\mathsf{fv}(\psi) \subset \mathsf{V}$, and let $\mathcal{M} = (\mathbb{A}, \sigma)$ be a structure. There is $\mathbb{A}'$ isomorphic to $\mathbb{A}$ and $\sigma' \in \mathcal{S}_{\mathsf{V}}$ such that $\mathcal{M} \models \psi$ just in case $(\mathbb{A}', \sigma') \models \psi$.*

Proof. Choose α such that σ differs from $\sigma' \in \mathcal{S}_{\mathsf{V}}$ by the bijective α, and let $\mathbb{A}' = \mathbb{A} \cdot \alpha$.

When $\mathcal{R}^*(\mathbb{A}_0, \mathbb{A}_\ell, \alpha)$, then $\mathcal{W}^*_X(\mathbb{A}_0, \mathbb{A}_\ell, \alpha)$ will mean that either $\mathbb{A}_0$ and $\mathbb{A}_\ell$ are isomorphic, or else $\alpha = \alpha_0 \circ \ldots \alpha_{\ell-1}$ and for each i from 0 to $\ell-1$, $\mathcal{W}_{X_i}(\mathbb{A}_i, \mathbb{A}_{i+1}, \alpha_i)$, where we let $X_0 = X$ and $X_{i+1} = X_i \cdot \alpha_i$. A set S of realizable skeletons is *X-irreducible* if whenever $\mathbb{A}, \mathbb{A}' \in S$ and $\mathcal{W}^*_X(\mathbb{A}, \mathbb{A}', \alpha)$, then $\mathbb{A} = \mathbb{A}'$.

Proposition 15. *If X, Y, Z are finite sets of atoms, there is an integer M such that the following holds. Whenever S is an X-irreducible set of realizable skeletons in which $\mathbb{A} \in S$ implies $\mathsf{non}_{\mathbb{A}} = Y$ and $\mathsf{unique}_{\mathbb{A}} = Z$, then $|S| \leq M$.*

Hence, for any finite Y, Z, there is a finite maximal X-irreducible set of realizable skeletons $\mathbb{A}$ with $\mathsf{non}_{\mathbb{A}} = Y, \mathsf{unique}_{\mathbb{A}} = Z$; we refer to this set as $\mathcal{I}_{X,Y,Z}$. $\mathcal{I}_{X,Y,Z}$ is computable.

Proof. As in the proof of Theorem 3, there are only finitely many strands no two of which are unifiable by α such that α is injective on X and bijective for $\mathsf{non}, \mathsf{unique}$. In particular, there are $i(j^{k+x+1})$, where i, j, k are as before and $x = |X|$. If the longest role of length ℓ, then there are at most $(i(j^{k+x+1}))^\ell$ choices which nodes to include in a skeleton, and thus only a finite number of choices of ordering.

In this way, we generate a sufficiently big set of skeletons. We may discard all skeletons from this set that are not realizable, since realizability is decidable (Proposition 6). We next discard any skeleton with an X-well-behaved reduction to an earlier skeleton. The result will be a maximal X-irreducible set. Hence, $\mathcal{I}_{X,Y,Z}$ is a computable set.

There is arbitrariness in the choice of $\mathcal{I}_{X,Y,Z}$, in the sense that there are different sets that fulfill the conditions. However, if $\mathcal{I}, \mathcal{I}'$ are two such sets, and $\mathbb{A} \in \mathcal{I}$ but $\mathbb{A} \notin \mathcal{I}'$, then $\mathbb{A}$ has an X-well-behaved reduction to some member of $\mathcal{I}'$.

Theorem 4. *Suppose that Π imposes no origination constraints. Satisfiability and validity are decidable for closed security goals*

$$\forall y_1, \ldots, y_n \,.\, H \supset C$$

where $H \in \mathcal{L}_\Pi$ is quantifier-free, and $C \in \mathcal{L}^0_\Pi$ is preserved under well-behaved reductions.

Proof. Suppose that $\forall y_1, \ldots, y_n \,.\, H \supset C$ is a formula of the form shown, and let $\psi = H \supset C$, where $\mathsf{V} = \mathsf{fv}(\psi) \subset \{y_1, \ldots, y_n\}$. Let NV be the set of variables x such that $\mathsf{nonp}(x)$ occurs at least once in H. Let UV be the set of variables x such that $\mathsf{uniquep}(x)$ occurs at least once in H. Let $\mathcal{S}_\mathsf{V}$ be a finite maximal set of variable assignments such that $\sigma, \sigma' \in \mathcal{S}_\mathsf{V}$ and $\sigma \sim_\mathsf{V} \sigma'$ implies $\sigma = \sigma'$, as in Proposition 13.

Choose any $\sigma \in \mathcal{S}_\mathsf{V}$. Let $X = \sigma(\mathsf{V})$, and choose any $Y \subset \sigma(\mathsf{NV})$ and $Z \subset \sigma(\mathsf{UV})$; there are only finitely many choices. For each choice, $\mathcal{I}_{X,Y,Z}$ as in Proposition 15 is finite. Thus, there are finitely many $\mathcal{M} = (\mathbb{A}, \sigma)$ with $\sigma \in \mathcal{S}_\mathsf{V}$ and $\mathbb{A} \in \mathcal{I}_{X,Y,Z}$ for some X, Y, Z constructed using σ. Call this finite set of structures $\mathfrak{A}$.

We would like to show that for any structure $(\mathbb{A}_0, \sigma_0)$, there is a structure $(\mathbb{A}_3, \sigma_3) \in \mathfrak{A}$ such that $(\mathbb{A}_3, \sigma_3) \models \psi$ if and only if $(\mathbb{A}_0, \sigma_0) \models \psi$. By Proposition 14, we may replace $(\mathbb{A}_0, \sigma_0)$ by $(\mathbb{A}_1, \sigma_1)$ with $\sigma_1 \in \mathcal{S}_\mathsf{V}$ and $(\mathbb{A}_0, \sigma_0) \models \psi$ if and only if $(\mathbb{A}_1, \sigma_1) \models \psi$. Define $Y_2 = \mathsf{non}_{\mathbb{A}_1} \cap \sigma_1(\mathsf{NV})$, $Z_2 = \mathsf{unique}_{\mathbb{A}_1} \cap \sigma_1(\mathsf{UV})$, $\sigma_2 = \sigma_1$, and

$$\mathbb{A}_2 = (\mathsf{N}_{\mathbb{A}_1}, \preceq_{\mathbb{A}_1}, Y_2, Z_2).$$

We must check that $(\mathbb{A}_2, \sigma_2) \models \psi$ if and only if $(\mathbb{A}_1, \sigma_1) \models \psi$. But ψ is $H \supset C$, and $(\mathbb{A}_2, \sigma_2) \models C$ if and only if $(\mathbb{A}_1, \sigma_1) \models C$, because $C \in \mathcal{L}^0_\Pi$. In fact, restricted to the vocabulary of $\mathcal{L}^0_\Pi$, $(\mathbb{A}_2, \sigma_2)$ is the same model as $(\mathbb{A}_1, \sigma_1)$; they differ only on unique and non, which are not expressed in $\mathcal{L}^0_\Pi$.

Since H is quantifier-free, its truth value in any model $(\mathbb{A}, \sigma)$ depends only on σ, not on the variants of σ. Whenever $\mathsf{nonp}(x)$ occurs in H, $(\mathbb{A}_2, \sigma_2) \models \mathsf{nonp}(x)$ if and only if $(\mathbb{A}_1, \sigma_1) \models \mathsf{nonp}(x)$, and likewise for $\mathsf{uniquep}(x)$. Thus, $(\mathbb{A}_2, \sigma_2) \models H$ if and only if $(\mathbb{A}_1, \sigma_1) \models H$.

So $(\mathbb{A}_2, \sigma_2)$ and $(\mathbb{A}_1, \sigma_1)$ agree on both C and H, and hence on ψ.

If $\mathbb{A}_2 \notin \mathcal{I}_{\sigma_2(\mathsf{V}),Y_2,Z_2}$, then by maximality, there is a $\sigma_2(\mathsf{V})$-preserving reduction of $\mathbb{A}_2$ to some member $\mathbb{A}_3 \in \mathcal{I}_{\sigma_2(\mathsf{V}),Y_2,Z_2}$. By assumption, this reduction leaves the truth-value of C unchanged. By Proposition 12, it leaves the truth value of H unchanged. Thus it leaves the truth value of ψ unchanged.

Proposition 16. *If $C \in \mathcal{L}^0_\Pi$ has the form $\exists x_1, \ldots, x_k \,.\, \phi^r_m(\boldsymbol{y})$, where the x_i may appear in $\boldsymbol{y}$, then C is preserved under well-behaved reductions.*

If $C \in \mathcal{L}^0_\Pi$ has the form $\forall x_1, \ldots, x_k \,.\, \neg(\phi^r_m(\boldsymbol{y}))$, where the x_i may appear in $\boldsymbol{y}$, then C is preserved under well-behaved reductions.

Proof. Suppose $\exists x_1, \ldots, x_k \,.\, \phi^r_m(\boldsymbol{y})$ is true in $(\mathbb{A}, \sigma)$. Let $\boldsymbol{a}^r$ be the atoms occurring in r. Then σ has an $\boldsymbol{x}$-variant σ' such that $\mathbb{A}$ contains a strand s such that $\mathsf{tr}(s) = \mathsf{tr}(r \cdot \alpha_0)$, where $a^r_i \cdot \alpha_0 = \sigma(x_i)$ for each i. A well-behaved reduction to $\mathbb{A}'$ via α maps s to some strand s' of $\mathbb{A}'$-height $\geq m$. The strand s' ensures that $\exists x_1, \ldots, x_k \,.\, \phi^r_m(\boldsymbol{y})$ is true in $(\mathbb{A}', \sigma \circ \alpha)$.

Conversely, if $\exists x_1, \ldots, x_k \,.\, \phi^r_m(\boldsymbol{y})$ is true in $(\mathbb{A}', \sigma \circ \alpha)$, there is such a strand s' in $\mathbb{A}$. By the injectiveness of α, any preimage of s' will be a strand s of role r with suitable parameters; at least one of them is of $\mathbb{A}$-height $\geq m$, as ϕ is surjective.

Universally quantified negative formulas are symmetrical.

Corollary 5. *If Π imposes no origination constraints, then authentication and non-disclosure goals for Π are decidable.*

It seems likely that a version of this result also containing ordering (recency) will also be true, with a version of Theorem 1 (and of the notion of reduction) that preserves some of the ordering of $\mathbb{A}'$.

8 Conclusion

We have studied *skeletons* and *preskeletons*, and the *homomorphisms* that relate them. Our main result is that a class of formulas of a first order language $\mathcal{L}_\Pi$ are decidable for protocols Π without origination assumptions. The bulk of concrete protocol analysis may be carried out using these formulas, or their enrichments containing ordering information.

The category of skeletons under homomorphisms is useful for other reasons [2]. It motivates a practical algorithm for protocol analysis, to find out just what can happen when a protocol is executed. This algorithm may be used whether Π makes origination assumptions or not, although it is not guaranteed to terminate in the former case. However, many protocols may be shown to have a single possible *shape* that all realizable skeletons share, and many others have a small finite number of shapes; protocol analysis may be automated by generating this set and observing what is true in it. This gives an efficient way to answer questions about protocols, unlike the one embedded in the proofs of Theorems 4.

Acknowledgment We are grateful to Iliano Cervesato and Dusko Pavlovic, who pointed out an error in an early version of this paper.

References

1. Bruno Blanchet and Andreas Podelski. Verification of cryptographic protocols: Tagging enforces termination. In Andrew D. Gordon, editor, *Foundations of Software Science and Computation Structures*, number 2620 in LNCS, pages 136–152. Springer, April 2003.
2. Shaddin Doghmi, Joshua Guttman, and F. Javier Thayer. The shapes of bundles. MTR 05 B 02, The MITRE Corp., 2004.
3. Nancy Durgin, Patrick Lincoln, John Mitchell, and Andre Scedrov. Multiset rewriting and the complexity of bounded security protocols. *Journal of Computer Security*, 12(2):247–311, 2004. Initial version appeared in *Workshop on Formal Methods and Security Protocols*, 1999.
4. Joshua D. Guttman. Security goals: Packet trajectories and strand spaces. In Roberto Gorrieri and Riccardo Focardi, editors, *Foundations of Security Analysis and Design*, volume 2171 of *LNCS*, pages 197–261. Springer Verlag, 2001.
5. Joshua D. Guttman and F. Javier Thayer. Authentication tests and the structure of bundles. *Theoretical Computer Science*, 283(2):333–380, June 2002.
6. James Heather and Steve Schneider. Toward automatic verification of authentication protocols on an unbounded network. In *Proceedings, 13th Computer Security Foundations Workshop*. IEEE Computer Society Press, July 2000.
7. James A. Heather and Steve A. Schneider. A decision procedure for the existence of a rank function. *Journal of Computer Security*, 2005. Forthcoming.
8. Jonathan K. Millen and Vitaly Shmatikov. Constraint solving for bounded-process cryptographic protocol analysis. In *8th ACM Conference on Computer and Communications Security (CCS '01)*, pages 166–175. ACM, 2001.
9. Roger Needham and Michael Schroeder. Using encryption for authentication in large networks of computers. *Communications of the ACM*, 21(12), December 1978.
10. R. Ramanujam and S. P. Suresh. Decidability of context-explicit security protocols. *Journal of Computer Security*, 13(1):135–166, 2005. Preliminary version appeared in WITS '03, *Workshop on Issues in the Theory of Security*, Warsaw, April 2003.

A Additional Strand Notions

Definition 15. Fix a strand space Σ:

1. The subterm relation $\sqsubset$ is the smallest reflexive, transitive relation such that $t \sqsubset \{|g|\}_K$ if $t \sqsubset g$, and $t \sqsubset g \,\hat{}\, h$ if either $a \sqsubset g$ or $a \sqsubset h$.
   (Hence, for $K \in \mathsf{K}$, we have $K \sqsubset \{|g|\}_K$ only if $K \sqsubset g$ already.)
2. There is an edge $n_1 \rightarrow n_2$ iff $\mathrm{term}(n_1) = +t$ or $+_c\, t$ and $\mathrm{term}(n_2) = -t$ or $-_a\, t$ for $t \in \mathsf{A}$. $n_1 \Rightarrow n_2$ means $n_1 = s \downarrow i$ and $n_2 = s \downarrow i+1 \in \mathcal{N}$.
   $n_1 \Rightarrow^* n_2$ (respectively, $n_1 \Rightarrow^+ n_2$) means that $n_1 = s \downarrow i$ and $n_2 = s \downarrow j \in \mathcal{N}$ for some s and $j \geq i$ (respectively, $j > i$).
3. Suppose I is a set of terms. The node $n \in \mathcal{N}$ is an *entry point* for I iff $\mathrm{term}(n) = +t$ for some $t \in I$, and whenever $n' \Rightarrow^+ n$, $\mathrm{term}(n') \notin I$. t *originates* on $n \in \mathcal{N}$ iff n is an entry point for $I = \{t' : t \sqsubset t'\}$.
4. A term t is *uniquely originating in* $S \subset \mathcal{N}$ iff there is a unique $n \in S$ such that t originates on n, and *non-originating* if there is no such $n \in S$.

Definition 16. A *penetrator strand* is an s where $\mathsf{tr}(s)$ is one of:

M_t: $\langle +a \rangle$ where a is a text, principal name, or nonce
K_K: $\langle +K \rangle$ where K is a key
$\mathsf{C}_{g,h}$: $\langle -g,\ -h,\ +(\mathsf{tag} \mathbin{\hat{}} g \mathbin{\hat{}} h) \rangle$
$\mathsf{S}_{g,h}$: $\langle -(\mathsf{tag} \mathbin{\hat{}} g \mathbin{\hat{}} h),\ +g,\ +h \rangle$
$\mathsf{E}_{h,K}$: $\langle -K,\ -h,\ +\{\!|h|\!\}_K \rangle$
$\mathsf{D}_{h,K}$: $\langle -K^{-1},\ -\{\!|h|\!\}_K,\ +h \rangle$

If s is a penetrator strand, $s \downarrow j$ is a *penetrator node*; otherwise it is *regular*.

Proposition 17. *If s is a penetrator strand of kind* M, K, C, *etc., and α is a replacement, then $s \cdot \alpha$ is a penetrator strand of the same kind,* M, K, C, *etc.*

Definition 17 (Bundles). Let $\mathcal{B} = \langle \mathcal{N}_\mathcal{B}, (\rightarrow_\mathcal{B} \cup \Rightarrow_\mathcal{B}) \rangle$ be a finite acyclic subgraph of $\langle \mathcal{N}, (\rightarrow \cup \Rightarrow) \rangle$. $\mathcal{B}$ is a *bundle* if:

1. If $n_2 \in \mathcal{N}_\mathcal{B}$ is negative, there is a unique n_1 such that $n_1 \rightarrow_\mathcal{B} n_2$.
2. If $n_2 \in \mathcal{N}_\mathcal{B}$ and $n_1 \Rightarrow n_2$ then $n_1 \Rightarrow_\mathcal{B} n_2$.

The *height* of a strand s in $\mathcal{B}$ is the largest i such that $s \downarrow i \in \mathcal{N}_\mathcal{B}$. The *bundle ordering* on $\mathcal{B}$ is the smallest reflexive, transitive relation $\preceq_\mathcal{B}$ such that $n_1 \rightarrow_\mathcal{B} n_2$ implies $n_1 \preceq_\mathcal{B} n_2$, and $n_1 \Rightarrow_\mathcal{B} n_2$ implies $n_1 \preceq_\mathcal{B} n_2$.

By acyclicity and finiteness, we have:

Proposition 18. *If $\mathcal{B}$ is a bundle, $\preceq_\mathcal{B}$ is a well-founded partial order.*