

On the Symbolic Analysis of Low-Level Cryptographic Primitives: Modular Exponentiation and the Diffie-Hellman Protocol *

Michele Boreale
Dipartimento di Sistemi e Informatica
Università di Firenze
Via Lombroso 6/17, 50137 Firenze, Italy
boreale@dsi.unifi.it

Maria Grazia Buscemi
Dipartimento di Informatica
Università di Pisa
Via F. Buonarroti 2, 56100 Pisa, Italy
buscemi@di.unipi.it

Abstract

Automatic methods developed so far for analysis of security protocols only model a limited set of cryptographic primitives (often, only encryption and concatenation) and abstract from low-level features of cryptographic algorithms. This paper is an attempt towards closing this gap. We propose a symbolic technique and a decision method for analysis of protocols based on modular exponentiation, such as Diffie-Hellman key exchange. We introduce a protocol description language along with its semantics. Then, we propose a notion of symbolic execution and, based on it, a verification method. We prove that the method is sound and complete with respect to the language semantics.

1 Introduction

During the last decade, a lot of research effort has been directed towards automatic analysis of crypto-protocols. Tools based on finite-state methods ([15, 17, 21]) take advantage of a well established model-checking technology, and are very effective at finding bugs. Infinite-state approaches, based on a variety of symbolic techniques ([3, 5, 9, 13, 19]), have emerged over the past few years. Implementations of these techniques (e.g. [24, 6, 23]) are still at an early stage. However, symbolic methods seem to be very promising in two respects. First, at least when the number of sessions is bounded, they can accomplish a complete exploration of the protocol's state space: thus they provide *proofs or disproofs* of correctness - under Dolev-Yao-like [12] assumptions

- even though the protocol's state space is infinite. Second, symbolic methods usually rely on representations of data which help to control very well state-explosion induced by communications.

The application of automatic methods has mostly been confined to protocols built around 'black-box' enciphering and hashing functions. In this paper, we take a step towards broadening the scope of symbolic techniques, so as to include a class of low-level cryptographic operations. In particular, building on the general framework proposed in [7], we devise a complete analysis method for protocols that depend on modular exponentiation operations, like the Diffie-Hellman key-exchange [11]. We expect that our methodology may be adapted to other low-level primitives (like RSA encryption).

The Diffie-Hellman protocol is intended for exchange of a secret key over an insecure medium, without prior sharing of any secret. The protocol has two public parameters: a large prime p and a generator α for the multiplicative group $\mathcal{Z}_p^* = \{1, \dots, p-1\}$. Assuming A and B want to establish a shared secret key, they proceed as follows. First, A generates a random private value $n_A \in \mathcal{Z}_p^*$ and B generates a random private value $n_B \in \mathcal{Z}_p^*$. Next, A and B exchange their public values ($\text{exp}(x, y)$ denotes $x^y \bmod p$):

1. $A \longrightarrow B : \text{exp}(\alpha, n_A)$
2. $B \longrightarrow A : \text{exp}(\alpha, n_B)$.

Finally, A computes the key as $K = \text{exp}(\text{exp}(\alpha, n_B), n_A) = \text{exp}(\alpha, n_A \times n_B)$, and B computes the key as $K = \text{exp}(\text{exp}(\alpha, n_A), n_B) = \text{exp}(\alpha, n_A \times n_B)$. Now A and B share the value K , and A can use it to, say, encrypt a secret datum d and

*This work has been partially supported by EU within the FET - Global Computing initiative, projects MIKADO and PROFUNDIS and by MIUR project NAPOLI.

send it to B :

$$3. \quad A \longrightarrow B : \{d\}_K.$$

The protocol’s security depends on the difficulty of the discrete logarithm problem: it is computationally infeasible to compute y if only x and $\exp(x, y)$ are known.

When defining a model for low-level protocols of this sort, one is faced with two conflicting requirements. On one hand, one should be accurate in accounting for the operations involved in the protocol (exponentiation, product) and their ‘relevant’ algebraic laws; even operations that are not explicitly mentioned in protocols, but that are considered feasible (like taking the k^{th} root modulo a prime, and division) must be accounted for, because an adversary could in principle take advantage of them.

On the other hand, one must be careful in keeping the model effectively analysable. In this respect, recent undecidability results on related problems of equational unification [14] indicate that some degree of abstraction is unavoidable. The limitations of our model are discussed in Section 2. Technically, we simplify the model by avoiding explicit commutativity laws and by keeping a free algebra model and ordinary unification. In fact, we ‘promote’ commutativity to non-determinism. As an example, upon evaluation of the expression $\exp(\exp(\alpha, n), m)$, an attacker will non-deterministically produce $\exp(\alpha, m \times n)$ or $\exp(\alpha, n \times m)$. The intuition is that if there is some action that depends on these two terms being equal modulo \times -commutativity, then there is an execution trace of the protocol where this action will take place. This seems reasonable in view of the fact that *we only consider safety properties* (i.e., ‘no bad action ever takes place’).

Here is a more precise description of our work. In Section 2, paralleling [7], we introduce a syntax for expressions (including $\exp(\cdot, \cdot)$ and related operations), along with a notion of evaluation. Based on this, we present a small protocol description language akin to the applied pi [1] and its (concrete) semantics. The latter assumes a Dolev-Yao adversary and is therefore infinitary. In Section 3, we introduce a finitary symbolic semantics, which relies on a form of narrowing strategy, and discuss its relationship with the concrete semantics. A verification method based on the symbolic semantics is presented in Section 4: the main result is Theorem 4.3, which asserts the correctness and completeness of the method with respect to the concrete model. It is remarkable that the presence of the modular root operation plays a crucial role in the completeness proof. In Section 5, we illustrate how the method discovers the well known man-in-the-middle attack on the Diffie-Hellman

protocol. A discussion on further research is in Section 6. A separate Appendix contains some most technical details. Due to space limitations, proofs are sketched or omitted but will appear in a forthcoming full version of the paper.

Related Work Very recent work by Millen and Shmatikov [20] shows how to reduce the symbolic analysis problem in the presence of modular exponentiation and multiplication plus encryption to the solution of quadratic Diophantine equations; decidability, however, remains an open issue. Pereira and Quisquater first [22] proposed a technique for analysing group Diffie-Hellman protocols in the presence of an attacker with restricted capabilities (e.g. no symmetric encryption), though not facing the issue of decidability. Blanchet’s model [4] abstracts away from operations like inverse, root extraction and random number generation. The resulting method may give rise to false attacks and may not terminate. Closely related to our problem is also protocol analysis in the presence of the xor operation, which has been recently proven to be decidable by Chevalier *et al.* [8] and, independently, by Comon-Lundh and Shmatikov [10].

2 The model

We recall the concept of *frame* from [7], and tailor it to the case of modular exponentiation and multiplication.

Expressions and messages We consider two countable disjoint sets of *names* $m, n, \dots \in \mathcal{N}$ and *variables* $x, y, \dots \in \mathcal{V}$. The set \mathcal{N} is in turn partitioned into a countable set of *local names* $a, b, \dots \in \mathcal{LN}$ and a countable set of *environmental names* $\underline{a}, \underline{b}, \dots \in \mathcal{EN}$: these two sets represent infinite supplies of fresh basic values (keys, random numbers, ...) at disposal of processes and of the (hostile) environment, respectively. It is technically convenient also to consider a set of *marked variables* $\hat{x}, \hat{y}, \hat{z}, \dots \in \hat{\mathcal{V}}$, which will be used as placeholders for generic messages known to the environment. The set $\mathcal{N} \cup \mathcal{V} \cup \hat{\mathcal{V}}$ is ranged over by u, v, \dots . Given a signature Σ of function symbols f, g, \dots , each coming with its arity (constants have arity 0), we denote by \mathcal{E}_Σ the algebra of terms (or *expressions*) on $\mathcal{N} \cup \mathcal{V} \cup \Sigma$, given by the grammar:

$$\zeta, \eta ::= u \mid f(\tilde{\zeta})$$

where $\tilde{\zeta}$ is a tuple of terms of the expected length. A *term context* $C[\cdot]$ is a term with a hole that can be filled with any expression ζ , thus yielding an expression $C[\zeta]$.

Definition 2.1 (a frame for exponentiation). A frame \mathcal{F} is a triple $(\Sigma, \mathcal{M}, \downarrow)$, where:

- Σ is a signature;
- $\mathcal{M} \subseteq \mathcal{E}_\Sigma$ is a set of messages M, N, \dots ;
- $\downarrow \subseteq \mathcal{E}_\Sigma \times \mathcal{E}_\Sigma$ is an evaluation relation.

A frame for exponentiation and shared key cryptography, $\mathcal{F}_{DH} = (\Sigma, \mathcal{M}, \downarrow)$, is presented in Table 1. We write $\zeta \downarrow \eta$ for $(\zeta, \eta) \in \downarrow$ and say that ζ evaluates to η .

Besides shared-key encryption $\{\zeta\}_\eta$ and decryption $\text{dec}_\eta(\zeta)$ (with η used as the key), the other symbols of Σ represent arithmetic operations modulo a fixed and public prime number, which is kept implicit. In particular, we have exponentiation $\text{exp}(\zeta, \eta)$, root extraction $\text{root}(\zeta, \eta)$, a constant α that represents a public generator and two constants for multiplicative unit (unit, 1), two distinct symbols for the product $\text{mult}(\zeta, \eta)$ and its result $\zeta \times \eta$, three symbols, $\text{inv}(\zeta)$, $\text{inv}'(\zeta)$ and ζ^{-1} , representing the multiplicative inverse operation. The reason for using different symbols for the same operation is discussed below. All the underlying operations are computationally feasible¹. A message is either a product of up to l values, for a fixed l , or a key or a message encrypted under a key. A key can be an atomic object, or an exponential with basis α and a product as exponent ($\text{exp}(\alpha, F)$).

Evaluation (\downarrow) is the reflexive and transitive closure of an auxiliary relation \rightsquigarrow , as presented in Table 1. There, we use $\zeta_1 \times \zeta_2 \times \dots \times \zeta_n$ as a shorthand for $\zeta_1 \times (\zeta_2 \times \dots \times \zeta_n)$, while (i_1, \dots, i_n) is any permutation of $(1, \dots, n)$. The relation \rightsquigarrow is terminating, but not confluent. In fact, the non-determinism of \rightsquigarrow is intended to model the commutativity and the associativity of the product operation, as reflected in the rule (MULT). Also note rule (ROOT): in modular arithmetic, taking the k^{th} root amounts to raising to the $(k^{-1} \bmod p-1)^{\text{th}}$ power. The adoption of distinct symbols for product (mult and \times), inverse (inv, inv' and $()^{-1}$), and unit (unit and 1), and the form of the rules, ensure termination of both \rightsquigarrow and of the induced narrowing relation, to be introduced in the next section.

The choice of the above message and rule formats corresponds to imposing the following restrictions on the attacker and on the honest participants:

¹Note that the inverse of k modulo n is defined only if $\text{gcd}(k, n) = 1$ (see [18]). Our model abstracts away from this condition. Another abstraction we make is that just one operation is used to model both inverse mod p and inverse mod $p-1$ (the latter operation arises only inside exponents).

1. there is a fixed upper bound (l) on the number of factors;
2. product and inverse operations cannot be applied to exponentials and to encrypted terms;
3. exponentiation starts from the basis α , and exponents can only be products.

More accurately, starting from a term obeying the above conditions, an attacker is capable of ‘deducing’ all - though not necessarily only - AC variants of the message represented by the term, in a sense made precise below. Terms not obeying these restrictions are just not guaranteed to produce any message. Restriction (1) might be relaxed at the cost of introducing a class of mult_l operations, for each $l \geq 0$, but for simplicity we shall stick to the above model in this paper.

Below, we define a deduction relation which expresses how the environment can generate new messages starting from an initial set of messages S . Note that environmental names and marked variables are treated as terms known to the environment. We denote by $\mathcal{P}_f(X)$ the set of finite subsets of X .

Definition 2.2 (deduction relation). For $S \subseteq \mathcal{M}$, the set $\mathcal{H}(S)$ is inductively defined by the following rules:

$$\begin{aligned} \mathcal{H}^0(S) &= S \cup \mathcal{EN} \cup \widehat{\mathcal{V}} \\ \mathcal{H}^{i+1}(S) &= \mathcal{H}^i(S) \cup \{f(\tilde{\zeta}) : f \in \Sigma, \tilde{\zeta} \subseteq \mathcal{H}^i(S)\} \\ \mathcal{H}(S) &= \bigcup_{i \geq 0} \mathcal{H}^i(S). \end{aligned}$$

The deduction relation $\vdash \subseteq \mathcal{P}_f(\mathcal{M}) \times \mathcal{M}$ is defined by:

$$S \vdash M \quad \text{iff} \quad \exists \zeta \in \mathcal{H}(S) : \zeta \downarrow M.$$

Example 2.3. Consider a set $S = \{n_A, \text{exp}(\alpha, n_B)\}$. Then, $S \vdash \text{exp}(\alpha, n_A \times n_B)$ and $S \vdash \text{exp}(\alpha, n_B \times n_A)$. A further example involves the use of the root extraction. Consider a set $S = \{\{m\}_{\text{exp}(\alpha, k \times l)}, \text{exp}(\alpha, k \times h), h, l\}$. Then, $S \vdash m$ since there exists $\zeta \in \mathcal{H}(S)$, $\zeta = \text{dec}_\eta(\{m\}_{\text{exp}(\alpha, k \times l)})$, with $\eta = \text{exp}(\text{root}(\text{exp}(\alpha, k \times h), h), l)$, s.t. $\zeta \downarrow m$.

Processes We consider a variant of the applied pi [1]. The syntax of *agents* is presented in Table 2.

Terms ζ and η range over \mathcal{E}_Σ . We consider a set \mathcal{L} of *labels* which is ranged over by a, b, \dots . A unique public channel is assumed, thus input and output labels (a, b, \dots) are simply ‘tags’ attached to process actions for ease of reference. The only construct (let) for evaluating expressions replaces the ad-hoc constructs found in the spi-calculus [2] for encryption, decryption and other cryptographic operations. The construct

SIGNATURE	$\Sigma = \{ \alpha, \text{unit}, 1, \{\cdot\}_{(\cdot)}, \text{dec}_{(\cdot)}(\cdot), \text{exp}(\cdot, \cdot), \text{root}(\cdot, \cdot), \cdot \times \cdot, \text{mult}(\cdot, \cdot), \text{inv}(\cdot), \text{inv}'(\cdot), (\cdot)^{-1} \}$
FACTORS	$f ::= u \mid u^{-1}$
PRODUCTS	$F ::= 1 \mid f_1 \times \dots \times f_k$
KEYS	$K, H ::= f \mid \text{exp}(\alpha, F)$
MESSAGES	$M, N ::= F \mid K \mid \{M\}_K$
(DEC)	$\text{dec}_\eta(\{\zeta\}_\eta) \rightsquigarrow \zeta$
(MULT)	$\text{mult}(\zeta_1 \times \dots \times \zeta_k, \zeta_{k+1} \times \dots \times \zeta_n) \rightsquigarrow \zeta_{i_1} \times \dots \times \zeta_{i_n} \quad 1 \leq k < n \leq l$
(INV ₁)	$\text{inv}(\zeta_1 \times \dots \times \zeta_n) \rightsquigarrow \text{inv}'(\zeta_1) \times \dots \times \text{inv}'(\zeta_n) \quad n \leq l$
(INV ₂)	$\text{inv}'(\zeta^{-1}) \rightsquigarrow \zeta \quad (\text{INV}_3) \quad \text{inv}'(\zeta) \rightsquigarrow \zeta^{-1} \quad (\text{INV}_4) \quad \text{inv}'(\zeta) \times \zeta \rightsquigarrow \text{unit}$
(UNIT ₁)	$\text{unit} \times \zeta \rightsquigarrow \zeta \quad (\text{UNIT}_2) \quad \text{unit} \rightsquigarrow 1$
(EXP)	$\text{exp}(\text{exp}(\xi, \eta), \zeta) \rightsquigarrow \text{exp}(\xi, \text{mult}(\eta, \zeta))$
(ROOT)	$\text{root}(\text{exp}(\xi, \eta), \zeta) \rightsquigarrow \text{exp}(\xi, \text{mult}(\eta, \text{inv}(\zeta)))$
(CTX)	$\frac{\zeta \rightsquigarrow \zeta'}{C[\zeta] \rightsquigarrow C[\zeta']}$
	EVALUATION $\zeta \downarrow \eta \quad \text{iff} \quad \zeta \rightsquigarrow^* \eta$

Table 1: \mathcal{F}_{DH} , a frame for modular exponentiation

$A, B ::=$	$\mathbf{0}$ (null)	$A = (\text{new } n_A) \overline{\mathbf{a1}}\langle \text{exp}(\alpha, n_A) \rangle. \mathbf{a2}(x).$
	$\mathbf{a}(x). A$ (input)	$\text{let } z = \text{exp}(x, n_A) \text{ in } \overline{\mathbf{a3}}\langle \{d\}_z \rangle. \mathbf{0}$
	$\overline{\mathbf{a}}\langle \zeta \rangle. A$ (output)	$B = (\text{new } n_B) \mathbf{b1}(y). \overline{\mathbf{b2}}\langle \text{exp}(\alpha, n_B) \rangle.$
	$\text{let } x = \zeta \text{ in } A$ (evaluation)	$\text{let } w = \text{exp}(y, n_B) \text{ in } \mathbf{b3}(t).$
	$[\zeta = \eta]A$ (matching)	$\text{let } t' = \text{dec}_w(t) \text{ in } B'$
	$A \parallel B$ (parallel composition)	$P = A \parallel B.$
	$(\text{new } a) A$ (restriction)	

Table 2: Syntax of agents

we just describe a one-session version of the protocol.

Here B' is a continuation of B after the reception of the encrypted datum d .

$(\text{new } a) A$ denotes the creation of a new name a , which is private to A . Note that the occurrences of variable x are bound in input and let operators.

Given the presence of binders for variables, notions of *free variables*, $\text{v}(A) \subseteq \mathcal{V}$, *substitution* $[\zeta/u]$ for any ζ and u , and *alpha-equivalence* arise as expected. We shall identify alpha-equivalent agents. We denote by $\text{en}(A)$ the set of environmental names occurring in A . An agent A is said to be *closed* or a *process* if $\text{v}(A) = \emptyset$; the set of processes \mathcal{P} is ranged over by P, Q, \dots

Example 2.4 (the Diffie-Hellman key exchange). *The process P below is a formalisation of the Diffie-Hellman protocol presented in the introduction. For simplicity,*

Operational Semantics The semantics of the calculus is given in terms of a transition relation \longrightarrow , which will sometimes be referred to as ‘concrete’ (as opposed to the ‘symbolic’ one we shall introduce later on). We model the state of the system as a pair $\langle s, P \rangle$, where s records the current environment’s knowledge (i.e., the sequence of messages the environment has ‘seen’ on the network up to a given moment) and P is a process term. Formally, an *action* is a term of the form $\mathbf{a}\langle M \rangle$ (*input action*) or $\overline{\mathbf{a}}\langle M \rangle$ (*output action*), for a label and M a message. The set of actions Act is ranged over by α, β, \dots , while the set Act^* is ranged over by s, s', \dots . String concatenation is written ‘.’. We denote by $\text{act}(s)$ and $\text{msg}(s)$ the set of actions and mes-

sages, respectively, appearing in s , and write ' $s \vdash M$ ' for $\text{msg}(s) \vdash M$.

Below we define *traces*, that is, sequences of actions that may result from the interaction between a process and its environment. In traces, each message received by a process (input message) can be synthesised from the knowledge the environment has previously acquired.

Definition 2.5 (traces and configurations). A trace is a string $s \in \text{Act}^*$ such that for each s_1, s_2 and $a \langle M \rangle$, if $s = s_1 \cdot a \langle M \rangle \cdot s_2$ then $s_1 \vdash M$. A configuration, written as $\langle s, P \rangle$, is a pair consisting of a ground trace s and a process P . A configuration is initial if $\text{en}(s, P) = \emptyset$. Configurations are ranged over by $\mathcal{C}, \mathcal{C}', \dots$.

The concrete transition relation on configurations is defined by the rules in Table 3. Each action taken by the process is recorded in the configuration's first component. Rule (INP) makes the transition relation infinitely-branching, as M ranges over the infinite set $\{M : s \vdash M\}$. In rule (OUT), ζ is evaluated and the action takes place only if the outcome is a message. By rule (LET), the evaluation of ζ replaces any free occurrence of y in P . In rule (NEW), a function $\text{new}_{\mathcal{L}\mathcal{N}}(V)$ is assumed that yields a local name $a \notin V$. No handshake communication is provided: all messages go through the environment (rule (PAR)). Note that, by (OUT), the non-determinism of \rightsquigarrow is lifted to processes, and this is used to render commutativity and associativity of product. We are of course ignoring any 'efficiency' consideration, given that the model is infinite anyway (rule (INP)).

Properties We focus on correspondence assertions of the kind 'for every generated trace, whenever action β occurs in the trace, then action α must have occurred at some previous point in the trace'. Formally, given a configuration $\langle s, P \rangle$ and a trace s' , we say that $\langle s, P \rangle$ *generates* s' , written $\langle s, P \rangle \searrow s'$, if $\langle s, P \rangle \xrightarrow{*} \langle s', P' \rangle$ for some P' . A substitution θ is a finite partial map from $\mathcal{V} \cup \widehat{\mathcal{V}}$ to the set of messages \mathcal{M} . For any object t (i.e. variable, message, process, trace, ...), we denote by $t\theta$ the result of simultaneously replacing each $x \in \text{v}(t) \cap \text{dom}(\theta)$ by $\theta(x)$. We let ρ range over ground substitutions, i.e. substitutions that map variables to ground messages.

Definition 2.6 (properties and satisfaction relation). Let α and β be actions and s be a trace. We say that α occurs prior to β in s if whenever $s = s' \cdot \beta \cdot s''$ then $\alpha \in \text{act}(s')$. For $\text{v}(\alpha) \subseteq \text{v}(\beta)$, we write $s \models \alpha \leftrightarrow \beta$, and say s satisfies $\alpha \leftrightarrow \beta$, if for each ground substitution ρ it holds that $\alpha\rho$ occurs prior to $\beta\rho$ in s . We say that a configuration \mathcal{C} satisfies $\alpha \leftrightarrow \beta$, and write $\mathcal{C} \models \alpha \leftrightarrow \beta$, if all traces generated by \mathcal{C} satisfy $\alpha \leftrightarrow \beta$.

Assertions $\alpha \leftrightarrow \beta$ can express interesting authentication (see [7]) and secrecy properties. In particular, secrecy in the style of [3] can be set within our framework by assuming a conventional 'absurd' action \perp that it is nowhere used in agent expressions. Thus, the formula $\perp \leftrightarrow \alpha$ means that action α should never take place.

Example 2.7 (Diffie-Hellman, continued). The property that the protocol P in Example 2.4 should not leak the datum d can be expressed also by saying that the adversary will never be capable of synthesising d , without prior knowledge of it. This can be formalised by extending P with a 'guardian' process $\mathbf{g}(t).\mathbf{0}$ that at any time can pick up one message from the network and then stop: $S = P \parallel \mathbf{g}(t).\mathbf{0}$. Then we check whether this guardian can ever pick d from the network. This means checking whether the initial configuration $\mathcal{C}_{DH} = \langle \epsilon, S \rangle$ ($\epsilon =$ empty trace) satisfies the property $\text{Secret}(d) = \perp \leftrightarrow \mathbf{g}(d)$, i.e. whether $\mathcal{C} \models \text{Secret}(d)$. Note that, by definition of deduction relation (Def. 2.2), $\epsilon \vdash \alpha$.

3 Symbolic Semantics

We equip the frame \mathcal{F}_{DH} with a *symbolic* evaluation relation (\downarrow_s), which is in agreement with its concrete counterpart (\downarrow). Intuitively, $\zeta \downarrow_\theta \eta$ means that ζ evaluates to η under all instances ρ of θ . The main advantage of the symbolic evaluation relation with respect to the concrete one is that infinitely many pairs (ζ, η) such that $\zeta \downarrow \eta$ can be represented by means of a single judgement $\zeta_0 \downarrow_\theta \eta_0$, for appropriate ζ_0, θ, η_0 .

Formally, let us denote by $\text{mgu}(t_1, t_2)$ a chosen *most general unifier* (mgu) of t_1 and t_2 , that is, a unifier θ of t_1 and t_2 such that any other unifier is a composition of substitutions θ and θ' , written $\theta\theta'$, for some θ' . Also, for t_1, t'_1, t_2, t'_2 terms, $\text{mgu}(t_1 = t'_1, t_2 = t'_2)$ stands for $\theta \text{mgu}(t_2\theta, t'_2\theta)$, where $\theta = \text{mgu}(t_1, t'_1)$, if such mgus exist. The symbolic evaluation relation \downarrow_s of \mathcal{F}_{DH} is presented in Table 4: it is defined as the reflexive and transitive closure of the relation $\rightsquigarrow_s^\theta$.

Lemma 3.1. Relation $\rightsquigarrow_s^\theta$ is image-finite and terminating. Hence, \downarrow_s is image-finite.

The notions of symbolic traces and symbolic configurations are defined below. The main difference from their concrete counterparts (Def. 2.5) is that no condition on input messages is required.

Definition 3.2 (symbolic traces and configurations). A symbolic trace is a string $s \in \text{Act}^*$ such that: (a) $\text{en}(s) = \emptyset$, and (b) for each s_1, s_2, α and x , if $s =$

(INP)	$\langle s, a(x).P \rangle \longrightarrow \langle s \cdot a\langle M \rangle, P[M/x] \rangle$	$s \vdash M$
(OUT)	$\langle s, \bar{a}\langle \zeta \rangle.P \rangle \longrightarrow \langle s \cdot \bar{a}\langle M \rangle, P \rangle$	$\zeta \downarrow M$
(LET)	$\langle s, \text{let } y = \zeta \text{ in } P \rangle \longrightarrow \langle s, P[\eta/y] \rangle$	$\zeta \downarrow \eta$
(MATCH)	$\langle s, [\zeta = \eta]P \rangle \longrightarrow \langle s, P \rangle$	$\zeta \downarrow \xi, \eta \downarrow \xi$
(NEW)	$\langle s, (\text{new } a)P \rangle \longrightarrow \langle s, P \rangle$	$a = \text{new}_{\mathcal{LN}}(V)$
(PAR)	$\frac{\langle s, P \rangle \longrightarrow \langle s', P' \rangle}{\langle s, P \parallel Q \rangle \longrightarrow \langle s', P' \parallel Q \rangle}$	

plus symmetric version of (PAR).

In (NEW), V is the set of free names in the source configuration.

Table 3: Rules for the concrete transition relation (\longrightarrow)

(DEC _s)	$\text{dec}_\eta(\zeta) \overset{\theta}{\rightsquigarrow}_s x_1 \theta$	$\theta = \text{mgu}(\zeta = \{x_1\}_{x_2}, \eta = x_2)$
(MULT _s)	$\text{mult}(\zeta_1, \zeta_n) \overset{\theta}{\rightsquigarrow}_s (x_{i_1} \times \dots \times x_{i_n}) \theta$	$\begin{cases} 1 \leq k < n \leq l, \\ \theta = \text{mgu}(\zeta_1 = x_1 \times \dots \times x_k, \\ \zeta_2 = x_{k+1} \times \dots \times x_n) \end{cases}$
(INV1 _s)	$\text{inv}(\zeta) \overset{\theta}{\rightsquigarrow}_s (\text{inv}'(x_{i_1}) \times \dots \times \text{inv}'(x_{i_n})) \theta$	$\begin{cases} 1 \leq n \leq l, \\ \theta = \text{mgu}(\zeta = x_1 \times \dots \times x_n) \end{cases}$
(INV2 _s)	$\text{inv}'(\zeta) \overset{\theta}{\rightsquigarrow}_s x_1 \theta$	$\theta = \text{mgu}(\zeta, x_1^{-1})$
(INV3 _s)	$\text{inv}'(\zeta) \overset{\epsilon}{\rightsquigarrow}_s \zeta^{-1}$	(INV4 _s) $\text{inv}'(\zeta) \times \eta \overset{\theta}{\rightsquigarrow}_s \text{unit}$ $\theta = \text{mgu}(\zeta, \eta)$
(UNIT1 _s)	$\text{unit} \times \zeta \overset{\epsilon}{\rightsquigarrow}_s \zeta$	(UNIT2 _s) $\text{unit} \overset{\epsilon}{\rightsquigarrow}_s 1$
(EXP1 _s)	$\text{exp}(x, \zeta) \overset{\theta}{\rightsquigarrow}_s \text{exp}(\alpha, \text{mult}(x_1, \zeta))$	$\theta = [\text{exp}(\alpha, x_1)/x]$
(EXP2 _s)	$\text{exp}(\text{exp}(\xi, \eta), \zeta) \overset{\epsilon}{\rightsquigarrow}_s \text{exp}(\xi, \text{mult}(\eta, \zeta))$	
(ROOT1 _s)	$\text{root}(x, \zeta) \overset{\theta}{\rightsquigarrow}_s \text{exp}(\alpha, \text{mult}(x_1, \text{inv}(\zeta)))$	$\theta = [\text{exp}(\alpha, x_1)/x]$
(ROOT2 _s)	$\text{root}(\text{exp}(\xi, \eta), \zeta) \overset{\epsilon}{\rightsquigarrow}_s \text{exp}(\xi, \text{mult}(\eta, \text{inv}(\zeta)))$	
(CTX _s)	$\frac{\zeta \overset{\theta}{\rightsquigarrow}_s \zeta'}{C[\zeta] \overset{\theta}{\rightsquigarrow}_s C\theta[\zeta']}$	

SYMBOLIC EVALUATION $\zeta \downarrow_\theta \eta$ iff $\zeta \overset{\theta_1}{\rightsquigarrow}_s \dots \overset{\theta_n}{\rightsquigarrow}_s \eta$ and $\theta = \theta_1 \dots \theta_n$

Variables x_1, \dots, x_n are fresh according to some arbitrary but fixed rule.

Table 4: Symbolic Evaluation Relation (\downarrow_s) for \mathcal{F}_{DH}

$s_1 \cdot \alpha \cdot s_2$ and $x \in v(\alpha) - v(s_1)$ then α is an input action. Symbolic traces are ranged over by σ, σ', \dots . A symbolic configuration, written $\langle \sigma, A \rangle_s$, is a pair composed by a symbolic trace σ and an agent A , such that $\text{en}(A) = \emptyset$ and $v(A) \subseteq v(\sigma)$.

The symbolic semantics is based on a symbolic transition relation \longrightarrow_s , defined in Table 5. There, a function $\text{new}_v(\cdot)$ is assumed such that, for any given $V \subseteq_{\text{fin}} \mathcal{V}$, $\text{new}_v(V)$ is a variable not in V . Moreover, $\mathcal{C} \xrightarrow{\theta}_s \mathcal{C}'$ stands for $\mathcal{C} \longrightarrow_s \mathcal{C}'$, where θ is the substitution applied to \mathcal{C}' in the reduction step.

Note that, unlike the concrete semantics, input variables are *not* instantiated immediately (rule (INP_s)). Rather, the input message is represented as a free variable and constraints on this variable are added as soon as they are needed, and recorded via mgu 's. This may occur due to rules (OUT_s) , (LET_s) and (MATCH_s) .

Example 3.3. The simple example below shows how, after a first step, variable x is instantiated to name b by a (MATCH_s) -reduction:

$$\begin{aligned} \langle \epsilon, a(x) \cdot [x = b]P \rangle_s &\longrightarrow_s \langle a\langle x \rangle, [x = b]P \rangle_s \\ &\longrightarrow_s \langle a\langle b \rangle, P[b/x] \rangle_s. \end{aligned}$$

As a more elaborate example, let $P = \bar{a}\langle k \rangle \cdot a(x)$. let $z = \text{root}(x, k)$ in P' . After an output action and an input action, the symbolical evaluation of $\text{root}(x, k)$ produces a global substitution $\theta = [\text{exp}(\alpha, x_1)/x]$ (x_1 fresh), to be applied to the whole configuration, and a local substitution $\theta' = [\text{exp}(\alpha, x_1 \times k_{-1})/z]$. I.e.

$$\langle \epsilon, P \rangle_s \xrightarrow{*}_s \langle \sigma\theta, P'\theta\theta' \rangle_s, \quad \text{with } \sigma = \bar{a}\langle k \rangle \cdot a\langle x \rangle.$$

Whenever $\langle \sigma, A \rangle_s \xrightarrow{*}_s \langle \sigma', A' \rangle_s$ for some A' , we say that $\langle \sigma, A \rangle_s$ symbolically generates σ' , and write $\langle \sigma, A \rangle_s \searrow_s \sigma'$. The relation \longrightarrow_s is finitely-branching since \downarrow_s is. Therefore, each configuration generates a finite number of symbolic traces. It is important to stress that many symbolic traces are in fact nonsense – sequences of actions that cannot be instantiated to any concrete trace. For instance, consider process $P = a(y) \cdot \text{let } x = \text{dec}_k(y) \text{ in } \bar{a}\langle x \rangle \cdot \mathbf{0}$. The initial configuration $\langle \epsilon, P \rangle_s$ symbolically generates the trace $a\langle \{x_0\}_k \rangle \cdot \bar{a}\langle x_0 \rangle$, which is inconsistent, because the environment cannot generate the value k in $\{x_0\}_k$ (i.e. $\epsilon \not\vdash k$, hence $\epsilon \not\vdash \{x_0\}_k$). The problem of detecting these inconsistent traces, that might give rise to ‘false positives’ when checking protocol properties, will be faced in the next section.

Theorem 3.5 below establishes a correspondence between the concrete and the symbolic transition relations.

It relies on the notion of consistency, defined below. Recall that marked variables are intended to carry messages known by the environment. For any \hat{x} and any trace σ , we denote by $\sigma \setminus \hat{x}$ the longest prefix of σ not containing \hat{x} .

Definition 3.4 (consistency). Let $\sigma \in \text{Act}^*$ and ρ be a ground substitution. We say that ρ satisfies σ if $\sigma\rho$ is a ground trace and, for each $\hat{x} \in v(\sigma)$, it holds that $(\sigma \setminus \hat{x})\rho \vdash \rho(\hat{x})$. In this case we say that $\sigma\rho$ is a solution of σ , and that σ is consistent.

Note that, because of the requirements on input messages, any trace (Def. 2.5) is obviously consistent: e.g., map all variables to some environmental names.

Theorem 3.5 (concrete vs. symbolic semantics). \mathcal{C} be an initial configuration and s be a ground trace. Then $\mathcal{C} \searrow_s s$ if and only if there exists σ such that $\mathcal{C} \searrow_s \sigma$ and s is a solution of σ .

4 A Verification Method

A crucial point of the method we present is checking consistency of symbolic traces. As mentioned in the previous section, a symbolic trace σ need not have solutions (ground instances that are traces). The next result allows us to check consistency.

Proposition 4.1. Let σ be a symbolic trace. Then there exists a finite set of traces $\text{Refinement}(\sigma)$, which are instances of σ and have the following property: for any s , s is a solution of σ if and only if s is a solution of some $\sigma' \in \text{Refinement}(\sigma)$.

Proposition 4.1 implies that σ is consistent if and only if $\text{Refinement}(\sigma) \neq \emptyset$. Roughly, the set $\text{Refinement}(\sigma)$ is computed by repeatedly unifying each input message in σ to terms that can be synthesized out of previous messages in σ . We refer to the Appendix for details. Here we shall content ourselves with giving a couple of examples; the second example shows the important role of the root extraction operation.

Example 4.2.

1. Consider $\sigma = \bar{c}\langle \{m\}_k \rangle \cdot \bar{c}\langle \{n\}_k \rangle \cdot c\langle \{x\}_k \rangle$. Then $\text{Refinement}(\sigma) = \{ \sigma[m/x], \sigma[n/x] \}$.
2. Consider $\sigma' = \bar{c}\langle h \rangle \cdot c\langle x \rangle \cdot \bar{c}\langle \text{exp}(\alpha, k \times h) \rangle \cdot \bar{c}\langle \{m\}_{\text{exp}(\alpha, k \times x)} \rangle \cdot c\langle m \rangle$. Clearly, σ' is consistent: e.g., map x to h . And, indeed, $\text{Refinement}(\sigma') = \{ \sigma'' \}$ where $\sigma'' = \sigma'[\hat{x}/x]$. It is notable that the root extraction operation, though not mentioned

(INP _s)	$\langle \sigma, a(x).A \rangle_s \longrightarrow_s \langle \sigma \cdot a\langle x \rangle, A \rangle_s$	
(OUT _s)	$\langle \sigma, \bar{a}\langle \zeta \rangle.A \rangle_s \longrightarrow_s \langle \sigma\theta \cdot \bar{a}\langle M \rangle, A\theta \rangle_s$	$\zeta \downarrow_\theta M$
(LET _s)	$\langle \sigma, \text{let } x = \zeta \text{ in } A \rangle_s \longrightarrow_s \langle \sigma\theta, A\theta[\xi/x] \rangle_s$	$\zeta \downarrow_\theta \xi$
(MATCH _s)	$\langle \sigma, [\zeta = \eta].A \rangle_s \longrightarrow_s \langle \sigma\theta, A\theta \rangle_s$	$\begin{cases} \zeta \downarrow_{\theta_1} \xi_1, \eta \theta_1 \downarrow_{\theta_2} \xi_2, \\ \theta_3 = \text{mgu}(\xi_1\theta_2, \xi_2), \\ \theta = \theta_1\theta_2\theta_3 \end{cases}$
(NEW _s)	$\langle \sigma, (\text{new } a).A \rangle_s \longrightarrow_s \langle \sigma, A \rangle_s$	$a = \text{new}_{\mathcal{LN}}(V)$
(PAR _s)	$\frac{\langle \sigma, A \rangle_s \longrightarrow_s \langle \sigma', A' \rangle_s}{\langle \sigma, A \parallel B \rangle_s \longrightarrow_s \langle \sigma', A' \parallel B' \rangle_s}$	

plus symmetric version of (PAR_s). In the above rules it is assumed that:

- (i) $x = \text{new}_V(V)$, being V the set of free variables in the source configuration;
- (ii) $\text{msg}(\sigma)\theta \subseteq \mathcal{M}$;
- (iii) in rule (PAR_s), $B' = B\theta$ where $\langle \sigma, A \rangle_s \xrightarrow{\theta}_s \langle \sigma', A' \rangle_s$.

Table 5: Rules for symbolic transition relation (\longrightarrow_s)

in σ'' , is essential to prove that σ'' is a trace. In fact, the environment is capable of learning m only by computing the key of the encrypted message as $\text{exp}(\text{root}(\text{exp}(\alpha, k \times h), h), \hat{x})$.

Let $\alpha \leftrightarrow \beta$ be a property and \mathcal{C} an initial configuration. The verification method $\mathbf{M}(\mathcal{C}, \alpha \leftrightarrow \beta)$, presented in Table 6, checks whether $\mathcal{C} \models \alpha \leftrightarrow \beta$ or not. Moreover, if the property is not satisfied, $\mathbf{M}(\mathcal{C}, \alpha \leftrightarrow \beta)$ computes a trace violating the property, that is, an attack on \mathcal{C} .

To understand how the method works, it is convenient to consider the simple case $\alpha = \perp$, i.e. the verification of $\mathcal{C} \models \perp \leftrightarrow \beta$. This means verifying that in the *concrete* semantics, no instance of action β is ever executed starting from \mathcal{C} . By the correspondence between symbolic and concrete semantics (Theorem 3.5), this amounts to checking that for each σ symbolically generated by \mathcal{C} , no solution of σ contains an instance of β . The method proceeds as follows. First, one checks whether there is a mgu θ of γ and β , for every action γ of σ . If, for every σ , such a θ does not exist, or if it exists but $\sigma\theta$ is not consistent (this means that the check $\exists \sigma' \in \mathbf{Refinement}(\sigma\theta)$ at step 5 fails), then the property holds true, otherwise it does not, and the trace σ' violating the property is reported.

The correctness of the method in the general case is stated by Theorem 4.3 below. Note that the method always terminates, given that \longrightarrow_s is finite-branching and, hence, $\mathbf{Mod}_{\mathcal{C}}$ is finite.

Theorem 4.3 (correctness and completeness). *Let \mathcal{C} be an initial configuration and α and β be actions with*

$$v(\alpha) \subseteq v(\beta).$$

- (1) *If $\mathbf{M}(\mathcal{C}, \alpha \leftrightarrow \beta)$ returns (No, σ') then $\mathcal{C} \not\models \alpha \leftrightarrow \beta$. In particular, for any injective ground substitution $\rho : v(\sigma') \rightarrow \mathcal{EN}$, $\mathcal{C} \searrow \sigma'\rho$ and $\sigma'\rho \not\models \alpha \leftrightarrow \beta$.*
- (2) *If $\mathcal{C} \not\models \alpha \leftrightarrow \beta$ then $\mathbf{M}(\mathcal{C}, \alpha \leftrightarrow \beta)$ returns (No, σ') and for any injective ground substitution $\rho : v(\sigma') \rightarrow \mathcal{EN}$, $\mathcal{C} \searrow \sigma'\rho$ and $\sigma'\rho \not\models \alpha \leftrightarrow \beta$.*

Remark 4.4. *In practice, rather than generating the whole set of symbolic traces at once (step 1) and then checking the property, it is more convenient to work ‘on-the-fly’ and comparing every last symbolic action γ taken by the configuration against action β of the property $\alpha \leftrightarrow \beta$. The refinement procedure $\mathbf{Refinement}(\cdot)$ is invoked only when β and γ are unifiable.*

5 Example: Symbolic Analysis of the Diffie-Hellman Protocol

In this section we apply the verification method described in Section 4 to analyse the Diffie-Hellman protocol. We adopt the formalisation of the protocol presented in Examples 2.4 and 2.7.

The Diffie-Hellman protocol is subject to attacks from active adversaries. In terms of our model, discovering an attack on the protocol amounts to finding a ground trace s such that the initial configuration

$\mathbf{M}(\mathcal{C}, \alpha \leftrightarrow \beta)$

1. compute $\mathbf{Mod}_{\mathcal{C}} = \{\sigma \mid \mathcal{C} \searrow_s \sigma\}$;
2. **foreach** $\sigma \in \mathbf{Mod}_{\mathcal{C}}$ **do**
3. **foreach** action γ in σ **do**
4. **if** $\exists \theta = \text{mgu}(\beta, \gamma)$ **and**
5. $\exists \sigma' \in \mathbf{Refinement}(\sigma\theta)$ **where** $\sigma' = \sigma\theta\theta'$ **and**
6. $\alpha\theta\theta'$ does not occur prior to $\beta\theta\theta'$ in σ'
7. **then return**(No, σ');
8. **return**(Yes);

Table 6: The verification method

$\mathcal{C}_{DH} = \langle \epsilon, S \rangle \searrow_s$ and $s \neq \text{Secret}(d) = \perp \leftrightarrow \mathbf{g}\langle d \rangle$.
And, indeed, such an s exists and is as follows:

$$\overline{\mathbf{a1}}\langle \exp(\alpha, n_A) \rangle \cdot \mathbf{a2}\langle \exp(\alpha, \underline{n}_I) \rangle \cdot \overline{\mathbf{a3}}\langle \{d\}_H \rangle \cdot \mathbf{g}\langle d \rangle,$$

where \underline{n}_I is any environmental name and $K = \exp(\alpha, \underline{n}_I \times n_A)$.

Intuitively, the above trace corresponds to an attack in which the environment intercepts $\exp(\alpha, n_A)$, generates a private name \underline{n}_I and handles $\exp(\alpha, \underline{n}_I)$ to A . Then, A computes $K = \exp(\alpha, \underline{n}_I \times n_A)$ and erroneously believes K is a shared key known by A and B . Finally, A sends over the network the secret datum d encrypted under the corrupted shared key K and, so, d is revealed to the environment.

Now, we have to show how this attack is discovered by our verification method. First, consider the following symbolic execution starting from \mathcal{C}_{DH} :

$$\begin{aligned} \mathcal{C}_{DH} &\longrightarrow_s \longrightarrow_s \\ &\longrightarrow_s \langle \overline{\mathbf{a1}}\langle \exp(\alpha, n_A) \rangle \cdot \mathbf{a2}\langle x \rangle, \\ &\quad \text{let } z = \exp(x, n_A) \text{ in } \overline{\mathbf{a3}}\langle \{d\}_z \rangle \cdot \mathbf{0} \parallel B \parallel \mathbf{g}\langle t \rangle \cdot \mathbf{0} \rangle_s \\ &\xrightarrow{\theta_0} \langle \overline{\mathbf{a1}}\langle \exp(\alpha, n_A) \rangle \cdot \mathbf{a2}\langle \exp(\alpha, x_0) \rangle, \\ &\quad \overline{\mathbf{a3}}\langle \{d\}_K \rangle \cdot \mathbf{0} \parallel B\theta_0\theta_1 \parallel \mathbf{g}\langle t \rangle \cdot \mathbf{0} \rangle_s \quad (*) \\ &\longrightarrow_s \langle \overline{\mathbf{a1}}\langle \exp(\alpha, n_A) \rangle \cdot \mathbf{a2}\langle \exp(\alpha, x_0) \rangle \cdot \overline{\mathbf{a3}}\langle \{d\}_K \rangle, \\ &\quad \mathbf{0} \parallel B\theta_0\theta_1 \parallel \mathbf{g}\langle t \rangle \cdot \mathbf{0} \rangle_s \\ &\longrightarrow_s \langle \overline{\mathbf{a1}}\langle \exp(\alpha, n_A) \rangle \cdot \mathbf{a2}\langle \exp(\alpha, x_0) \rangle \cdot \overline{\mathbf{a3}}\langle \{d\}_K \rangle \cdot \mathbf{g}\langle t \rangle, \\ &\quad \mathbf{0} \parallel B\theta_0\theta_1 \parallel \mathbf{0} \rangle_s \\ &= \langle \sigma, \mathbf{0} \parallel B\theta_0\theta_1 \parallel \mathbf{0} \rangle_s, \end{aligned}$$

where $K = \exp(\alpha, x_0 \times n_A)$. In step (*), rule (LET_s) is applied, with $\exp(x, n_A) \downarrow_{\theta_0} \exp(\alpha, x_0 \times n_A)$, $\theta_0 = [\exp(\alpha, x_0)/x]$ (x_0 fresh), and $\theta_1 = [\exp(\alpha, x_0)/z]$.

Now, we execute the method step by step:

1. The symbolic model $\mathbf{Mod}_{\mathcal{C}_{DH}}$ is computed (in practice, symbolic traces are generated and checked ‘on-the-fly’).
2. The symbolic trace σ defined above is considered.

3,4. Action $\gamma = \mathbf{g}\langle t \rangle$ is found that unifies with $\beta = \mathbf{g}\langle d \rangle$, via $\theta = [d/t]$.

5. The set $\mathbf{Refinement}(\sigma\theta) = \{\sigma'\}$ is computed where $\sigma' = \sigma\theta\theta'$, and $\theta' = [\hat{x}_0/x_0]$. As stated by Theorem 4.1, σ' is a consistent trace. Note, in particular, that if we let $\sigma' = \sigma'' \cdot \mathbf{g}\langle d \rangle$, then $\sigma'' \vdash d$. Indeed, there exists $\zeta = \text{dec}_{\xi}(\{d\}_{\eta}) \in \mathcal{H}(\sigma'')$, with $\eta = \exp(\alpha, \hat{x}_0 \times n_A)$, $\xi = \exp(\exp(\alpha, n_A), \hat{x}_0)$, and $\zeta \rightsquigarrow_s \rightsquigarrow_s \text{dec}_{\eta}(\{d\}_{\eta}) \rightsquigarrow_s d$.

6. Action \perp does not appear in σ' , hence,

7. (No, σ') is returned. Note that s is retrieved as $\sigma'[\underline{n}_I/\hat{x}_0]$.

Of course, here we have proceeded by hand, but the symbolic traces generated by \mathcal{C}_{DH} are of the order of hundreds, hence we expect that an implementation of the method would detect this attack in a fraction of second.

It is straightforward to extend the frame of Section 2 with operations such as pairing, public key encryption and decryption, hashing and digital signature; we refer the reader to [7] for a detailed treatment of those operations. In particular, digital signature allows one to describe the authenticated version of the Diffie-Hellman protocol, where the public values exchanged are signed by the two partners, so that the man-in-the-middle attack is thwarted.

6 Conclusions and future work

We have presented a model for the analysis of protocols built around shared-key encryption and modular exponentiation. The model is less precise than others (notably [20]), its main limitation being a bound on the number of factors that may appear in any exponent. However, for such restricted model, we offer a decision

method. The model and the method smoothly carry over when including other common enciphering, signing and hashing primitives. We also believe the method is effective in practice, because the symbolic model is compact, and the refinement procedure at its heart is only invoked on demand and on single symbolic traces. We are in the process of integrating the technique presented here into the STA analysis tool ([6, 23]).

Our technical development has been confined to multiplication and exponentiation, but the methodology presented suggests directions for extensions to other low-level primitives. It certainly seems promising for the case of RSA encryption, which is based on exponentiation too.

Acknowledgements. The authors wish to thank the anonymous referees for their helpful comments.

References

- [1] M. Abadi, C. Fournet. Mobile Values, New Names, and Secure Communication. In *Conf. Rec. of POPL'01*, 2001.
- [2] M. Abadi, A.D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148(1):1-70, 1999.
- [3] R.M. Amadio, S. Lugiez. On the reachability problem in cryptographic protocols. In *Proc. of Concur'00*, LNCS 1877, Springer-Verlag, 2000. Full version: RR 3915, INRIA Sophia Antipolis.
- [4] B. Blanchet. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *Proc. of 14th Computer Security Foundations Workshop*, IEEE Computer Society Press, 2001.
- [5] M. Boreale. Symbolic Trace Analysis of Cryptographic Protocols. In *Proc. of ICALP'01*, LNCS 2076, Springer-Verlag, 2001.
- [6] M. Boreale, M. Buscemi. Experimenting with STA, a Tool for Automatic Analysis of Security Protocols. In *Proc. of SAC'02*, ACM Press, 2002.
- [7] M. Boreale and M. Buscemi. A Framework for the Analysis of Security Protocol. In *Proc. of CONCUR '02*, LNCS 2421. Springer-Verlag, 2002.
- [8] Y. Chevalier, R. Kuesters, M. Rusinowitch, and M. Turuani. An NP Decision Procedure for Protocol Insecurity with Xor. In *Proc. of LICS '03*, IEEE Computer Society Press, 2003.
- [9] H. Comon, V. Cortier, J. Mitchell. Tree automata with one memory, set constraints and ping-pong protocols. In *Proc. of ICALP'01*, LNCS 2076, Springer-Verlag, 2001.
- [10] H. Comon-Lundh and V. Shmatikov. Intruder Deductions, Constraint Solving and Insecurity Decision in Presence of Exclusive or. In *Proc. LICS '03*, IEEE Computer Society Press, 2003.
- [11] W. Diffie, M. Hellman. New Directions in Cryptography. *IEEE Transactions on Information Theory*, 22(6):644-654, 1976.
- [12] D. Dolev, A. Yao. On the security of public-key protocols. *IEEE Transactions on Information Theory*, 29(2):198-208, 1983.
- [13] A. Huima. Efficient infinite-state analysis of security protocols. In *Proc. of Workshop on Formal Methods and Security Protocols*, Trento, 1999.
- [14] D. Kapur, P. Narendran, and L. Wang. An E-unification Algorithm for Analyzing Protocols that Use Modular Exponentiation. In *Proc. of RTA '03*, LNCS 2706, Springer-Verlag, 2003.
- [15] G. Lowe. Breaking and Fixing the Needham-Schroeder Public-Key Protocol Using FDR. In *Proc. of TACAS'96*, LNCS 1055, Springer-Verlag, 1996.
- [16] G. Lowe. A Hierarchy of Authentication Specifications. In *Proc. of 10th IEEE Computer Security Foundations Workshop*, IEEE Computer Society Press, 1997.
- [17] W. Marrero, E.M. Clarke, S. Jha. Model checking for security protocols. Technical Report TR-CMU-CS-97-139, Carnegie Mellon University, 1997.
- [18] A Menezes, P.C. van Oorschot, and S. Vanstone. *Handbook of Applied Cryptography*, CRC Press, Boca Raton, 1997.
- [19] J. Millen, V. Shmatikov. Constraint solving for bounded-process cryptographic protocol analysis. In *Proc. of 8th ACM Conference on Computer and Communication Security*, ACM Press, 2001.
- [20] J. Millen and V. Shmatikov. Symbolic Protocol Analysis with Products and Diffie-Hellman Exponentiation. In *Proc. of 16th IEEE Computer Security Foundations Workshop*, IEEE Computer Society Press, 2003.

- [21] J.C. Mitchell, M. Mitchell, U Stern. Automated Analysis of Cryptographic Protocols Using Mur ϕ . In *Proc. of Symp. Security and Privacy*, IEEE Computer Society Press, 1997.
- [22] O. Pereira and J.J Quisquater. A Security Analysis of the Cliques Protocols Suites. In *Proc. of the 14th IEEE Computer Security Foundations Workshop*, IEEE Computer Society Press, 2001.
- [23] STA: a tool for trace analysis of cryptographic protocols. ML object code and examples. Available at <http://www.dsi.unifi.it/~boreale/tool.html>.
- [24] V. Vanack`ere. The TRUST Protocol Analyser, Automatic and Efficient Verification of Cryptographic Protocols. In *Proc. of Verify '02*, 2002.

A Computing Refinement

Here we move within the general framework of [7]². First, we introduce the concept of *basis*, which intuitively consists of the ‘building blocks’ of all messages deducible from a given σ . Given generic products F and G , we write $G \subset F$ to indicate that factors of G are strictly included in those of F .

Definition A.1 (a basis for \mathcal{F}_{DH}). *For each symbolic trace σ :*

$$\mathbf{b}_{DH}(\sigma) = \{ M \mid (\sigma \vdash M) \text{ and } (M \in \mathcal{LN} \cup \{\alpha, 1\} \\ \text{or } (M = F \text{ and } \sigma \not\vdash G, \forall G \subset F) \\ \text{or } (M = \exp(\alpha, F) \text{ and } \sigma \not\vdash G, \forall G \subset F) \\ \text{or } (M = \{M\}_K \text{ and } \sigma \not\vdash K)) \}.$$

In practice, for a given σ , the set $\mathbf{b}_{DH}(\sigma)$ can be effectively computed by an iterative procedure which repeatedly applies destructors (dec, root, mult, inv, inv’) to messages in σ , until some fixed point is reached. This procedure always terminates and yields a finite set of messages.

Definition A.2 (Refinement(σ)). *We let refinement, written \succ , be the least binary relation over symbolic traces given by the two rules below. In (REF₁), it is assumed: that σ' is the longest prefix of σ which is a trace, that $\sigma = \sigma' \cdot \mathbf{a}\langle M \rangle \cdot \sigma''$, for some σ'' , that $N, N' \notin \mathcal{V} \cup \widehat{\mathcal{V}}$,*

²With a small change in terminology: in [7], consistent symbolic traces are called ‘solved forms’, while **Refinement**(\cdot) is denoted by **SF**(\cdot)

and that $\theta \neq \epsilon$.

$$\begin{array}{c} \text{(REF}_1\text{)} \frac{N' \in \mathbf{b}(\sigma') \quad M = C[N] \quad \theta = \text{mgu}(N, N')}{\sigma \succ \sigma\theta\theta_0} \\ \text{(REF}_2\text{)} \frac{x \in \mathbf{v}(M)}{\sigma \succ \sigma[\hat{x}/x]} \end{array}$$

where $\theta_0 = [x/\hat{x} \mid \hat{x} \in \mathbf{v}(\sigma) \text{ and } |(\sigma\theta)\setminus\hat{x}| < |\sigma\setminus\hat{x}|]$.

For any symbolic trace σ , we let **Refinement**(σ) = $\{ \sigma' \mid \sigma \succ^* \sigma' \text{ and } \sigma' \text{ is a trace} \}$.

Rule (REF₁) implements the basic step of refinement: the subterm N of M gets unified, via θ , with an element of $\mathbf{b}(\sigma')$. By rule (REF₂) a variable x can become marked: this amounts to constraining the possible values of x to be messages known by the environment. (For technical reasons, marked variables sometimes need to be ‘unmarked’ back to plain variables, and this is achieved in (REF₁) via the renaming θ_0 .)

Lemma A.3. *Refinement(σ) can be effectively computed, and is finite.*

Proof. It follows by two facts: (a) \succ is an image-finite relation, and (b) infinite sequences of refinement steps cannot arise. As to the latter point, note that, since each (REF₁)-step eliminates at least one variable, any sequence of refinement steps can contain only finitely many (REF₁)-steps. After the last such step, rule (REF₂) can only be applied a finite number of times. \square