

YAPA: A Generic Tool for Computing Intruder Knowledge^{*}

Mathieu Baudet¹, Véronique Cortier², and Stéphanie Delaune³

¹ DCSSI, France

² LORIA, CNRS & INRIA project Cassis, France

³ LSV, ENS Cachan & CNRS & INRIA, France

Abstract. Reasoning about the knowledge of an attacker is a necessary step in many formal analyses of security protocols. In the framework of the applied pi calculus, as in similar languages based on equational logics, knowledge is typically expressed by two relations: deducibility and static equivalence. Several decision procedures have been proposed for these relations under a variety of equational theories. However, each theory has its particular algorithm, and none has been implemented so far.

We provide a generic procedure for deducibility and static equivalence that takes as input any convergent rewrite system. We show that our algorithm covers all the existing decision procedures for convergent theories. We also provide an efficient implementation, and compare it briefly with the more general tool ProVerif.

1 Introduction

Understanding security protocols often requires reasoning about the information accessible to an online attacker. Accordingly, many formal approaches to security rely on a notion of *deducibility* [18,19] that models whether a piece of data, typically a secret, is retrievable from a finite set of messages. Deducibility, however, does not always suffice to reflect the knowledge of an attacker. Consider for instance a protocol sending an encrypted Boolean value, say, a vote in an electronic voting protocol. Rather than deducibility, the key idea to express confidentiality of the plaintext is that an attacker should not be able to *distinguish* between the sequences of messages corresponding to each possible value.

In the framework of the applied pi-calculus [3], as in similar languages based on equational logics [10], indistinguishability corresponds to a relation called *static equivalence*: roughly, two sequences of messages are *statically equivalent* when they satisfy the same algebraic relations from the attacker's point of view. Static equivalence plays an important role in the study of guessing attacks (e.g. [1,5,13]), as well as for anonymity properties and electronic voting protocols (e.g. [17]). In several cases, this notion has also been shown to imply the more complex and precise notion of cryptographic indistinguishability [1,8], related to probabilistic polynomial-time Turing machines.

^{*} Kindly supported by ANR-07-SESU-002 AVOTÉ and ARA SSIA FormaCrypt.

We emphasize that both deducibility and static equivalence apply to observations on finite sets of messages, and do not take into account the dynamic behavior of protocols. Nevertheless, deducibility is used as a subroutine by many general decision procedures [12,11]. Besides, it has been shown that observational equivalence in the applied pi-calculus coincides with labeled bisimulation [3], that is, corresponds to checking a number of static equivalences and some standard bisimulation conditions.

Deducibility and static equivalence rely on an underlying equational theory for axiomatizing the properties of cryptographic functions. Many decision procedures [2,14] have been proposed to compute these relations under a variety of equational theories, including symmetric and asymmetric encryptions, signatures, exclusive OR, and homomorphic operators. However, except for the class of subterm convergent theories [2], which covers the standard flavors of encryption and signature, each of these decision results introduces a new procedure, devoted to a particular theory. Even in the case of the general decidability criterion given in [2], we note that the algorithm underlying the proof has to be adapted for each theory, depending on how the criterion is fulfilled.

Perhaps as a consequence of this fact, none of these decision procedures has been implemented so far. Up to our knowledge the only tool able to verify static equivalence is ProVerif [9,10]. This general tool can handle various equational theories and analyze security protocols under active adversaries. However termination of the verifier is not guaranteed in general, and protocols are subject to (safe) approximations.

The present work aims to fill this gap between theory and implementation and propose an efficient tool for deciding deducibility and static equivalence in a uniform way. It is initially inspired from a procedure for solving more general constraint systems related to active adversaries and equivalence of finite processes, presented in [5], with corrected extended version in [6] (in French). However, due to the complexity of the constraint systems, this decision procedure was only studied for subterm convergent theories, and remains too complex to enable an efficient implementation.

Our first contribution is to provide and study a generic procedure for checking deducibility and static equivalence, taking as input any convergent theory (that is, any equational theory described by a finite convergent rewrite system). We prove the algorithm sound and complete, up to explicit failure cases. Note that (unfailing) termination cannot be guaranteed in general since the problem of checking deducibility and static equivalence is undecidable, even for convergent theories [2]. To address this issue and turn our algorithm into a decision procedure for a given convergent theory, we provide two criteria. First, we define a syntactic criterion on the rewrite rules that ensures that the algorithm never fails. This criterion is enjoyed in particular by any convergent subterm theory, as well as the theories of blind signature and homomorphic encryption. Termination often follows from a simple analysis of the rules of the algorithm: as a proof of concept, we obtain a new decidability result for deducibility and static equivalence for the prefix theory, representing encryption in CBC mode.

Second, we provide a termination criterion based on deducibility: provided that failure cannot occur, termination on a given input is equivalent to the existence of some natural finite representation of deducible terms. As a consequence, we obtain that our algorithm can decide deducibility and static equivalence for all the convergent theories previously known to be decidable [2].

Our second contribution is an efficient implementation of this generic procedure, called YAPA. After describing the main features of the implementation, we report several experiments suggesting that our tool computes static equivalence faster and for more convergent theories than the general tool ProVerif [9,10]. Because of space constraints, proofs are in an extended version of this paper [7].

2 Preliminaries

2.1 Term Algebra

We start by introducing the necessary notions to describe cryptographic messages in a symbolical way. For modeling cryptographic primitives, we assume a given set of *function symbols* \mathcal{F} together with an arity function $\text{ar} : \mathcal{F} \rightarrow \mathbb{N}$. Symbols in \mathcal{F} of arity 0 are called *constants*. We consider a set of *variables* \mathcal{X} and a set of additional constants \mathcal{W} called *parameters*. The (usual, first-order) term algebra generated by \mathcal{F} over \mathcal{W} and \mathcal{X} is written $\mathcal{F}[\mathcal{W} \cup \mathcal{X}]$ with elements denoted by $T, U, T_1 \dots$. More generally, we write $\mathcal{F}'[A]$ for the least set of terms containing a set A and stable by application of symbols in $\mathcal{F}' \subseteq \mathcal{F}$.

We write $\text{var}(T)$ (resp. $\text{par}(T)$) for the set of variables (resp. parameters) that occur in a term T . These notations are extended to tuples and sets of terms in the usual way. The set of positions (resp. subterms) of a term T is written $\text{pos}(T) \subseteq \mathbb{N}^*$ (resp. $\text{st}(T)$). The subterm of T at position $p \in \text{pos}(T)$ is written $T|_p$. The term obtained by replacing $T|_p$ with a term U in T is denoted $T[U]_p$.

A (*finite, partial*) *substitution* σ is a mapping from a finite subset of variables, called its *domain* and written $\text{dom}(\sigma)$, to terms. The *image* of a substitution is its image as a mapping $\text{im}(\sigma) = \{\sigma(x) \mid x \in \text{dom}(\sigma)\}$. Substitutions are extended to endomorphisms of $\mathcal{F}[\mathcal{X} \cup \mathcal{W}]$ as usual. We use a postfix notation for their application. A term T (resp. a substitution σ) is *ground* iff $\text{var}(T) = \emptyset$ (resp. $\text{var}(\text{im}(\sigma)) = \emptyset$).

For our cryptographic purposes, it is useful to distinguish a subset \mathcal{F}_{pub} of \mathcal{F} , made of *public function symbols*, that is, intuitively, the symbols made available to the attacker. A *recipe* (or *second-order term*) $M, N, M_1 \dots$ is a term in $\mathcal{F}_{\text{pub}}[\mathcal{W} \cup \mathcal{X}]$, that is, a term containing no *private* (non-public) function symbols. A *plain term* (or *first-order term*) $t, r, s, t_1 \dots$ is a term in $\mathcal{F}[\mathcal{X}]$, that is, containing no parameters. A (*public, ground, non-necessarily linear*) *n-ary context* C is a recipe in $\mathcal{F}_{\text{pub}}[\mathbf{w}_1, \dots, \mathbf{w}_n]$, where we assume a fixed countable subset of parameters $\{\mathbf{w}_1, \dots, \mathbf{w}_n, \dots\} \subseteq \mathcal{W}$. If C is a *n-ary context*, $C[T_1, \dots, T_n]$ denotes the term obtained by replacing each occurrence of \mathbf{w}_i with T_i in C .

2.2 Rewriting

A *rewrite system* \mathcal{R} is a finite set of *rewrite rules* $l \rightarrow r$ where $l, r \in \mathcal{F}[\mathcal{X}]$ and $\text{var}(r) \subseteq \text{var}(l)$. A term S *rewrites* to T by \mathcal{R} , denoted $S \rightarrow_{\mathcal{R}} T$, if there exist $l \rightarrow r$ in \mathcal{R} , $p \in \text{pos}(S)$ and a substitution σ such that $S|_p = l\sigma$ and $T = S[r\sigma]_p$. We write $\rightarrow_{\mathcal{R}}^+$ for the transitive closure of $\rightarrow_{\mathcal{R}}$, $\rightarrow_{\mathcal{R}}^*$ for its reflexive and transitive closure, and $=_{\mathcal{R}}$ for its reflexive, symmetric and transitive closure.

A rewrite system \mathcal{R} is *convergent* if is *terminating*, i.e. there is no infinite chains $T_1 \rightarrow_{\mathcal{R}} T_2 \rightarrow_{\mathcal{R}} \dots$, and *confluent*, i.e. for every terms S, T such that $S =_{\mathcal{R}} T$, there exists U such that $S \rightarrow_{\mathcal{R}}^* U$ and $T \rightarrow_{\mathcal{R}}^* U$.

A term T is \mathcal{R} -*reduced* if there is no term S such that $T \rightarrow_{\mathcal{R}} S$. If $T \rightarrow_{\mathcal{R}}^* S$ and S is \mathcal{R} -reduced then S is a \mathcal{R} -*reduced form* of T . When this reduced form is unique (in particular if \mathcal{R} is convergent), we write $S = T \downarrow_{\mathcal{R}}$.

2.3 Equational Theories

We equip the signature \mathcal{F} with an equational theory represented by a set of equations \mathcal{E} of the form $s = t$ with $s, t \in \mathcal{F}[\mathcal{X}]$. The equational theory \mathbf{E} generated by \mathcal{E} is the least set of equations containing \mathcal{E} that is stable under the axioms of congruence (reflexivity, symmetry, transitivity, application of function symbols) and under application of substitutions. We write $=_{\mathbf{E}}$ for the corresponding relation on terms. Equational theories have proved very useful for modeling algebraic properties of cryptographic primitives [2,15].

We are particularly interested in theories \mathbf{E} that can be represented by a convergent rewrite system \mathcal{R} , i.e. theories for which there exists a convergent rewrite system \mathcal{R} such that the two relations $=_{\mathcal{R}}$ and $=_{\mathbf{E}}$ coincide. The rewrite system \mathcal{R} —and by extension the equational theory \mathbf{E} — is *subterm convergent* if, in addition, we have that for every rule $l \rightarrow r \in \mathcal{R}$, r is either a subterm of l or a ground \mathcal{R} -reduced term. This class encompasses the one of the same name used in [2], the class of *dwindling theories* used in [4], and the class of *public-collapsing theories* introduced in [16].

Example 1. Consider the signature $\mathcal{F}_{\text{enc}} = \{\text{dec}, \text{enc}, \langle -, - \rangle, \pi_1, \pi_2\}$. The symbols dec , enc and $\langle -, - \rangle$ are functional symbols of arity 2 that represent respectively the decryption, encryption and pairing functions, whereas π_1 and π_2 are functional symbols of arity 1 that represent the projection function on the first and the second component of a pair, respectively. The equational theory of pairing and symmetric (deterministic) encryption, denoted by \mathbf{E}_{enc} , is generated by the equations $\mathcal{E}_{\text{enc}} = \{\text{dec}(\text{enc}(x, y), y) = x, \pi_1(\langle x, y \rangle) = x, \pi_2(\langle x, y \rangle) = y\}$.

Motivated by the modeling of the ECB mode of encryption, we may also consider an encryption symbol that is homomorphic with respect to pairing:

$$\mathcal{E}_{\text{hom}} = \mathcal{E}_{\text{enc}} \cup \left\{ \begin{array}{l} \text{enc}(\langle x, y \rangle, z) = \langle \text{enc}(x, z), \text{enc}(y, z) \rangle \\ \text{dec}(\langle x, y \rangle, z) = \langle \text{dec}(x, z), \text{dec}(y, z) \rangle \end{array} \right\}.$$

If we orient the equations from left to right, we obtain two rewrite systems \mathcal{R}_{enc} and \mathcal{R}_{hom} . Both rewrite systems are convergent, only \mathcal{R}_{enc} is subterm convergent.

From now on, we assume a given equational theory E represented by a convergent rewrite system \mathcal{R} . A symbol f is *free* if f does not occur in \mathcal{R} . In order to model (an unbounded number of) random values possibly generated by the attacker, we assume that \mathcal{F}_{pub} contains infinitely many free public constants. We will use free private constants to model secrets, for instance the secret keys used to encrypt a message. Private (resp. public) free constants are closely related to bound (resp. free) *names* in the framework of the applied pi calculus [3]. Our formalism also allows one to consider non-constant private symbols.

3 Deducibility and Static Equivalence

In order to describe the cryptographic messages observed or inferred by an attacker, we introduce the following notions of deduction facts and frames.

A *deduction fact* is a pair, written $M \triangleright t$, made of a recipe $M \in \mathcal{F}_{\text{pub}}[\mathcal{W} \cup \mathcal{X}]$ and a plain term $t \in \mathcal{F}[\mathcal{X}]$. Such a deduction fact is *ground* if $\text{var}(M, t) = \emptyset$. A *frame*, denoted by letters $\varphi, \Phi, \Phi_0, \dots$, is a finite set of ground deduction facts. The *image* of a frame is defined by $\text{im}(\Phi) = \{t \mid M \triangleright t \in \Phi\}$. A frame Φ is *one-to-one* if $M_1 \triangleright t, M_2 \triangleright t \in \Phi$ implies $M_1 = M_2$.

A frame φ is *initial* if it is of the form $\varphi = \{w_1 \triangleright t_1, \dots, w_\ell \triangleright t_\ell\}$ for some distinct parameters $w_1, \dots, w_\ell \in \mathcal{W}$. Initial frames are closely related to the notion of frames in the applied pi-calculus [3]. The parameters w_i can be seen as labels that refer to the messages observed by an attacker. Given such an initial frame φ , we denote by $\text{dom}(\varphi)$ its *domain* $\text{dom}(\varphi) = \{w_1, \dots, w_\ell\}$. If $\text{par}(M) \subseteq \text{dom}(\varphi)$, we write $M\varphi$ for the term obtained by replacing each w_i by t_i in M . If in addition M is ground then $t = M\varphi$ is a ground plain term.

3.1 Deducibility, Recipes

Classically (see e.g. [2]), a ground term t is *deducible* modulo E from an initial frame φ if there exists $M \in \mathcal{F}_{\text{pub}}[\text{dom}(\varphi)]$ such that $M\varphi =_E t$. This corresponds to the intuition that the attacker may compute (infer) t from φ . For the purpose of our study, we generalize this notion to arbitrary frames, and even sets of (non-necessarily ground) deduction facts ϕ , using the notations \triangleright_ϕ and \triangleright_ϕ^E .

Definition 1 (deducibility). *Let ϕ be finite set of deductions facts, for instance a frame. We say that M is a recipe of t in ϕ , written $M \triangleright_\phi t$, iff there exist a (public, ground, non-necessarily linear) n -ary context C and some deduction facts $M_1 \triangleright t_1, \dots, M_n \triangleright t_n$ in ϕ such that $M = C[M_1, \dots, M_n]$ and $t = C[t_1, \dots, t_n]$. In that case, we say that t is syntactically deducible from ϕ , also written $\phi \vdash t$.*

We say that M is a recipe of t in ϕ modulo E , written $M \triangleright_\phi^E t$, iff there exists a term t' such that $M \triangleright_\phi t'$ and $t' =_E t$. In that case, we say that t is deducible from ϕ modulo E , written $\phi \vdash_E t$.

We note that $M \triangleright_\varphi t$ is equivalent to $M\varphi = t$ when φ is an initial frame and when t (or equivalently M) is ground.

Example 2. Consider the equational theory E_{enc} given in Example 1. Let $\varphi = \{w_1 \triangleright \langle \text{enc}(s_1, k), \text{enc}(s_2, k) \rangle, w_2 \triangleright k\}$ where s_1, s_2 and k are private constant symbols. We have that $\langle w_2, w_2 \rangle \triangleright_{\varphi} \langle k, k \rangle$, and $\text{dec}(\text{proj}_1(w_1), w_2) \triangleright_{\varphi}^{E_{\text{enc}}} s_1$.

3.2 Static Equivalence, Visible Equations

Deducibility does not always suffice for expressing the knowledge of an attacker. In particular, it does not account for the partial information that an attacker may obtain about secrets. This issue motivates the study of visible equations and static equivalence [3], defined as follows.

Definition 2 (static equivalence). *Let φ be an initial frame. The set of visible equations of φ modulo E is defined as*

$$\text{eq}_E(\varphi) = \{M \bowtie N \mid M, N \in \mathcal{F}_{\text{pub}}[\text{dom}(\varphi)], M\varphi =_E N\varphi\}$$

where \bowtie is a dedicated commutative symbol. Two initial frames φ_1 and φ_2 with the same domain are statically equivalent modulo E , written $\varphi_1 \approx_E \varphi_2$, if their sets of visible equations are equal, i.e. $\text{eq}_E(\varphi_1) = \text{eq}_E(\varphi_2)$.

This definition is in line with static equivalence in the applied pi calculus [3]. For the purpose of finitely describing the set of visible equations $\text{eq}_E(\varphi)$ of an initial frame, we introduce *quantified equations* of the form $\forall z_1, \dots, z_q. M \bowtie N$ where $z_1, \dots, z_q \in \mathcal{X}$, $q \geq 0$ and $\text{var}(M, N) \subseteq \{z_1, \dots, z_q\}$. In the following, finite sets of quantified equations are denoted Ψ, Ψ_0, \dots . We write $\Psi \models M \bowtie N$ when the ground equation $M \bowtie N$ is a consequence of Ψ in the usual, first-order logics with equality axioms for the relation \bowtie (that is, reflexivity, symmetry, transitivity and compatibility with symbols in \mathcal{F}_{pub}). When no confusion arises, we may refer to quantified equations simply as *equations*. As usual, quantified equations are considered up to renaming of bound variables.

Example 3. Consider again the equational theory E_{enc} given in Example 1. Let $\varphi_1 = \{w_1 \triangleright \text{enc}(c_0, k), w_2 \triangleright k\}$ and $\varphi_2 = \{w_1 \triangleright \text{enc}(c_1, k), w_2 \triangleright k\}$ where c_0, c_1 are public constants and k is a private constant. Let $\Psi_1 = \{\text{enc}(c_0, w_2) \bowtie w_1\}$ and $\Psi_2 = \{\text{enc}(c_1, w_2) \bowtie w_1\}$. We have that $\Psi_i \models \text{eq}_{E_{\text{enc}}}(\varphi_i)$ for $i = 1, 2$. Hence, $\text{eq}_{E_{\text{enc}}}(\varphi_1) \neq \text{eq}_{E_{\text{enc}}}(\varphi_2)$ and the two frames φ_1 and φ_2 are not statically equivalent. However, it can be shown that $\{w_1 \triangleright \text{enc}(c_0, k)\} \approx_{E_{\text{enc}}} \{w_1 \triangleright \text{enc}(c_1, k)\}$.

4 Main Procedure

In this section, we describe our algorithms for checking deducibility and static equivalence on convergent rewrite systems. After some additional notations, we present the core of the procedure, which consists of a set of transformation rules used to saturate a frame and a finite set of quantified equations. We then show how to use this procedure to decide deducibility and static equivalence, provided that saturation succeeds.

Soundness and completeness of the saturation procedure are detailed in Section 5. We provide sufficient conditions on the rewrite systems to ensure success of saturation in Section 6.

4.1 Decompositions of Rewrite Rules

Before stating the procedure, we introduce the following notion of *decomposition* to account for the possible superpositions of an attacker's context with a left-hand side of rewrite rule.

Definition 3 (decomposition). *Let n, p, q be non-negative integers. A (n, p, q) -decomposition of a term l (and by an extension of any rewrite rule $l \rightarrow r$) is a (public, ground, non-necessarily linear) context $D \in \mathcal{F}_{\text{pub}}[\mathcal{W}]$ such that $\text{par}(D) = \{\mathbf{w}_1, \dots, \mathbf{w}_{n+p+q}\}$ and $l = D[l_1, \dots, l_n, y_1, \dots, y_p, z_1, \dots, z_q]$ where*

- l_1, \dots, l_n are mutually-distinct non-variable terms,
- y_1, \dots, y_p and z_1, \dots, z_q are mutually-distinct variables, and
- $y_1, \dots, y_p \in \text{var}(l_1, \dots, l_n)$ whereas $z_1, \dots, z_q \notin \text{var}(l_1, \dots, l_n)$.

A decomposition D is proper if it is not a parameter (i.e. $D \neq \mathbf{w}_1$).

Example 4. Consider the rewrite rule $\text{dec}(\text{enc}(x, y), y) \rightarrow x$. This rule admits two proper decompositions up to permutation of parameters:

- $D_1 = \text{dec}(\text{enc}(\mathbf{w}_1, \mathbf{w}_2), \mathbf{w}_2)$ where $n = 0, p = 0, q = 2, z_1 = x, z_2 = y$;
- $D_2 = \text{dec}(\mathbf{w}_1, \mathbf{w}_2)$ where $n = 1, p = 1, q = 0, l_1 = \text{enc}(x, y)$ and $y_1 = y$.

4.2 Transformation Rules

To check deducibility and static equivalence, we proceed by saturating an initial frame, adding some deduction facts and equations satisfied by the frame. We consider *states* that are either the failure state \perp or a couple (Φ, Ψ) formed by a one-to-one frame Φ in \mathcal{R} -reduced form and a finite set of quantified equations Ψ .

Given an initial frame φ , our procedure starts from an initial state associated to φ , denoted by $\text{Init}(\varphi)$, obtained by reducing φ and replacing duplicated terms by equations. Formally, $\text{Init}(\varphi)$ is the result of a procedure recursively defined as follows: $\text{Init}(\emptyset) = (\emptyset, \emptyset)$, and assuming $\text{Init}(\varphi) = (\Phi, \Psi)$, we have

$$\text{Init}(\varphi \uplus \{w \triangleright t\}) = \begin{cases} (\Phi, \Psi \cup \{w \bowtie w'\}) & \text{if there exists some } w' \triangleright t \downarrow_{\mathcal{R}} \in \Phi \\ (\Phi \cup \{w \triangleright t \downarrow_{\mathcal{R}}\}, \Psi) & \text{otherwise.} \end{cases}$$

The main part of our procedure consists in saturating a state (Φ, Ψ) by means of the transformation rules described in Figure 1. The **A** rules are designed for applying a rewrite step on top of existing deduction facts. If the resulting term is already syntactically deducible then a corresponding equation is added (rule **A.1**); or else if it is ground, the corresponding deduction fact is added to the state (rule **A.2**); otherwise, the procedure may fail (rule **A.3**). The **B** rules are meant to add syntactically deducible subterms (rule **B.2**) or related equations (rule **B.1**). For technical reasons, rule **A.1** is parametrized by a function Ctx with values of the form M or \perp , and satisfying the following properties:

- (a) if $\phi \vdash t \downarrow_{\mathcal{R}}$, then for any Ψ and α , $\text{Ctx}(\phi \vdash_{\mathcal{R}}^? t, \Psi, \alpha) \neq \perp$;

A. Inferring deduction facts and equations by context reduction

Assume that

$$\begin{aligned} l &= D[l_1, \dots, l_n, y_1, \dots, y_p, z_1, \dots, z_q] \text{ is a proper decomposition of } (l \rightarrow r) \in \mathcal{R} \\ M_1 \triangleright t_1, \dots, M_{n+p} \triangleright t_{n+p} &\in \Phi \\ (l_1, \dots, l_n, y_1, \dots, y_p) \sigma &= (t_1, \dots, t_{n+p}) \end{aligned}$$

1. If there exists $M = \text{Ctx}(\Phi \cup \{z_1 \triangleright z_1, \dots, z_q \triangleright z_q\} \vdash_{\mathcal{R}}^? r\sigma, \Psi, (l, r, D, \sigma))$, then

$$(\Phi, \Psi) \Longrightarrow (\Phi, \Psi \cup \{\forall z_1, \dots, z_q. D[M_1, \dots, M_{n+p}, z_1 \dots, z_q] \bowtie M\}) \quad (\mathbf{A.1})$$

2. Else, if $(r\sigma) \downarrow_{\mathcal{R}}$ is ground, then

$$\begin{aligned} (\Phi, \Psi) \Longrightarrow & (\Phi \cup \{M_0 \triangleright (r\sigma) \downarrow_{\mathcal{R}}\}, \\ & \Psi \cup \{\forall z_1, \dots, z_q. D[M_1, \dots, M_{n+p}, z_1 \dots, z_q] \bowtie M_0\}) \end{aligned} \quad (\mathbf{A.2})$$

where $M_0 = D[M_1, \dots, M_{n+p}, \mathbf{a}, \dots, \mathbf{a}]$ for some fixed public constant \mathbf{a} .

3. Otherwise, $(\Phi, \Psi) \Longrightarrow \perp$ ($\mathbf{A.3}$)

B. Inferring deduction facts and equations syntactically

Assume that $M_0 \triangleright t_0, \dots, M_n \triangleright t_n \in \Phi \quad t = f(t_1, \dots, t_n) \in \text{st}(t_0) \quad f \in \mathcal{F}_{\text{pub}}$

1. If there exists M such that $(M \triangleright t) \in \Phi$,

$$(\Phi, \Psi) \Longrightarrow (\Phi, \Psi \cup \{f(M_1, \dots, M_n) \bowtie M\}) \quad (\mathbf{B.1})$$

2. Otherwise, $(\Phi, \Psi) \Longrightarrow (\Phi \cup \{f(M_1, \dots, M_n) \triangleright t\}, \Psi)$ ($\mathbf{B.2}$)

Fig. 1. Transformation rules

- (b) if $M = \text{Ctx}(\phi \vdash_{\mathcal{R}}^? t, \Psi, \alpha)$ then there exist M' and s such that $\Psi \models M \bowtie M'$, $M' \triangleright_{\phi} s$ and $t \rightarrow_{\mathcal{R}}^* s$. (This justifies the notation $\phi \vdash_{\mathcal{R}}^? t$ used to denote a specific deducibility problem.)

Note that a simple choice for $\text{Ctx}(\phi \vdash_{\mathcal{R}}^? t, \Psi, \alpha)$ is to solve the deducibility problem $\phi \vdash_{\mathcal{R}}^? t \downarrow_{\mathcal{R}}$ in the empty equational theory, and then return a corresponding recipe M , if any. (This problem is easily solved by induction on $t \downarrow_{\mathcal{R}}$.) Yet, optimizing the function Ctx is a nontrivial task: on the one hand, letting $\text{Ctx}(\phi \vdash_{\mathcal{R}}^? t, \Psi, \alpha) \neq \perp$ for more values ϕ, t, Ψ, α makes the procedure more likely to succeed; on the other hand, it is computationally more demanding. We explain in Section 6.1 the choice of Ctx made in our implementation.

We write \Longrightarrow^* for the transitive and reflexive closure of \Longrightarrow . The definitions of Ctx and of the transformation rules ensure that whenever $S \Longrightarrow^* S'$ and S is a state, then S' is also a state, with the same parameters unless $S' = \perp$.

Example 5. Consider the frame φ_1 previously described in Example 3. We can apply rule **A.1** as follows. Consider the rewrite rule $\text{dec}(\text{enc}(x, y), y) \rightarrow x$, the decomposition D_2 given in Example 4 and $t_1 = \text{enc}(c_0, k)$. We have $\text{Init}(\varphi_1) = (\varphi_1, \emptyset) \Longrightarrow (\varphi_1, \{\text{dec}(w_1, w_2) \bowtie c_0\})$. In other words, since we know the key k through w_2 , we can check that the decryption of w_1 by w_2 leads to the public

constant c_0 . Next we apply rule **B.1** as follows: $(\varphi_1, \{\text{dec}(w_1, w_2) \bowtie c_0\}) \Longrightarrow (\varphi_1, \{\text{dec}(w_1, w_2) \bowtie c_0, \text{enc}(c_0, w_2) \bowtie w_1\})$. No more rules can then modify the state.

Main theorem. We now state the soundness and the completeness of the transformation rules provided that a *saturated state* is reached, that is, a state $S \neq \perp$ such that $S \Longrightarrow S'$ implies $S' = S$. The technical lemmas involved in the proof are detailed in Section 5.

Theorem 1 (soundness and completeness). *Let \mathbb{E} be an equational theory generated by a convergent rewrite system \mathcal{R} . Let φ be an initial frame and (Φ, Ψ) be a saturated state such that $\text{Init}(\varphi) \Longrightarrow^* (\Phi, \Psi)$.*

1. For all $M \in \mathcal{F}_{\text{pub}}[\text{par}(\varphi)]$ and $t \in \mathcal{F}[\emptyset]$, we have

$$M\varphi =_{\mathbb{E}} t \quad \Leftrightarrow \quad \exists N, \Psi \models M \bowtie N \text{ and } N \triangleright_{\Phi} t \downarrow_{\mathcal{R}}$$

2. For all $M, N \in \mathcal{F}_{\text{pub}}[\text{par}(\varphi) \cup \mathcal{X}]$, we have that $M\varphi =_{\mathbb{E}} N\varphi \Leftrightarrow \Psi \models M \bowtie N$.

While the saturation procedure is sound and complete, it may not terminate, or *fail* if rule **A.3** becomes the only applicable rule. In Section 6, we explore several sufficient conditions to prevent failure and ensure termination.

4.3 Application to Deduction and Static Equivalence

Decision procedures for deduction and static equivalence follow from Theorem 1.

Algorithm for deduction. Let φ be an initial frame and t be a ground term. The procedure for checking $\varphi \vdash_{\mathbb{E}} t$ runs as follows:

1. Apply the transformation rules to obtain (if any) a saturated state (Φ, Ψ) such that $\text{Init}(\varphi) \Longrightarrow^* (\Phi, \Psi)$;
2. Return *yes* if there exists N such that $N \triangleright_{\Phi} t \downarrow_{\mathcal{R}}$ (that is, the \mathcal{R} -reduced form of t is syntactically deducible from Φ); otherwise return *no*.

Algorithm for static equivalence. Let φ_1 and φ_2 be two initial frames. The procedure for checking $\varphi_1 \approx_{\mathbb{E}} \varphi_2$ runs as follows:

1. Apply the transformation rules to obtain (if possible) two saturated states (Φ_1, Ψ_1) and (Φ_2, Ψ_2) such that $\text{Init}(\varphi_i) \Longrightarrow^* (\Phi_i, \Psi_i)$, $i = 1, 2$;
2. For $\{i, j\} = \{1, 2\}$, for every equation $(\forall z_1, \dots, z_\ell. M \bowtie N)$ in Ψ_i , check that $M\varphi_j =_{\mathbb{E}} N\varphi_j$ — that is, in other words, $(M\varphi_j) \downarrow_{\mathcal{R}} = (N\varphi_j) \downarrow_{\mathcal{R}}$;
3. If so return *yes*; otherwise return *no*.

5 Soundness and Completeness of the Saturation

The proof of Theorem 1 is based on three main lemmas. First, the transformation rules are sound in the sense that, along the saturation process, we add only deducible terms and valid equations with respect to the initial frame.

Lemma 1 (soundness). *Let φ be an initial frame and (Φ, Ψ) be a state such that $\text{Init}(\varphi) \Longrightarrow^* (\Phi, \Psi)$. Then, we have that*

1. $M \triangleright_{\Phi} t \Rightarrow M\varphi =_{\mathbf{E}} t$ for all $M \in \mathcal{F}_{\text{pub}}[\text{dom}(\varphi)]$ and $t \in \mathcal{F}[\emptyset]$;
2. $\Psi \models M \bowtie N \Rightarrow M\varphi =_{\mathbf{E}} N\varphi$ for all $M, N \in \mathcal{F}_{\text{pub}}[\text{dom}(\varphi) \cup \mathcal{X}]$.

The next two lemmas are dedicated to the completeness of **B** and **A** rules, respectively. Lemma 2 ensures that saturated states account for all the syntactic equations possibly visible. Lemma 3 deals with the reduction of a deducible term along the rewrite system \mathcal{R} . Using that \mathcal{R} is convergent, this allows us to prove that every deducible term from a saturated frame is syntactically deducible.

Lemma 2 (completeness, syntactic equations). *Let (Φ, Ψ) be a state, and M, N be two terms such that $M \triangleright_{\Phi} t$ and $N \triangleright_{\Phi} t$ for some term t . Then there exists (Φ', Ψ') such that $(\Phi, \Psi) \Longrightarrow^* (\Phi', \Psi')$ using **B** rules and $\Psi' \models M \bowtie N$.*

Lemma 3 (completeness, context reduction). *Let (Φ, Ψ) be a state and M, t, t' be three terms such that $M \triangleright_{\Phi} t$ and $t \rightarrow_{\mathcal{R}} t'$. Then, either $(\Phi, \Psi) \Longrightarrow^* \perp$ or there exist (Φ', Ψ') , M' and t'' such that $(\Phi, \Psi) \Longrightarrow^* (\Phi', \Psi')$, $M' \triangleright_{\Phi'} t''$ with $t' \rightarrow_{\mathcal{R}}^* t''$, and $\Psi' \models M \bowtie M'$.*

*Besides, in both cases, the corresponding derivation from (Φ, Ψ) can be chosen to consist of a number of **B** rules, possibly followed by one instance of **A** rule involving the same rewrite rule $l \rightarrow r$ as the rewrite step $t \rightarrow_{\mathcal{R}} t'$.*

6 Termination and Non-failure

In the previous section, we proved that saturated frames yield sound and complete characterizations of deducible terms and visible equations of their initial frames. Yet, the saturation procedure may still not terminate, or fail due to rule **A.3**. In this section, we study different conditions on the rewrite system \mathcal{R} so that failure never happens and/or termination is ensured.

6.1 A Syntactic Criterion to Prevent Failure

Our first criterion is syntactic and ensures that the algorithm never fails. It is enjoyed by a large class of equational theories, called *layered convergent*.

Definition 4 (layered rewrite system). *A rewrite system \mathcal{R} , and by extension its equational theory \mathbf{E} , are layered if there exists an ascending chain of subsets $\emptyset = \mathcal{R}_0 \subseteq \mathcal{R}_1 \subseteq \dots \subseteq \mathcal{R}_{N+1} = \mathcal{R}$ ($N \geq 0$), such that for every $0 \leq i \leq N$, for every rule $l \rightarrow r$ in $\mathcal{R}_{i+1} - \mathcal{R}_i$, for every (n, p, q) -decomposition $l = D[l_1, \dots, l_n, y_1, \dots, y_p, z_1, \dots, z_q]$, one of the following two conditions holds:*

- (i) $\text{var}(r) \subseteq \text{var}(l_1, \dots, l_n)$;
- (ii) there exist C_0, C_1, \dots, C_k and s_1, \dots, s_k such that
 - $r = C_0[s_1, \dots, s_k]$;
 - for each $1 \leq i \leq k$, $C_i[l_1, \dots, l_n, y_1, \dots, y_p, z_1, \dots, z_q]$ rewrites to s_i in zero or one step of rewrite rule in head position along \mathcal{R}_i .

In the latter case, we say that the context $C = C_0[C_1, \dots, C_k]$ is associated to the decomposition D of $l \rightarrow r$. Note that $C[l_1, \dots, l_n, y_1, \dots, y_p, z_1, \dots, z_q] \rightarrow_{\mathcal{R}_i}^* r$.

Proposition 1. *Assume that the function Ctx in use is maximal: for every ϕ and t , if there exists s such that $\phi \vdash s$ and $t \rightarrow_{\mathcal{R}}^* s$, then for any Ψ, α , $\text{Ctx}(\phi \vdash_{\mathcal{R}}^? t, \Psi, \alpha) \neq \perp$. Then, provided that \mathcal{R} is layered convergent, there exists no state (Φ, Ψ) from which $(\Phi, \Psi) \Longrightarrow \perp$ is the only applicable derivation.*

Practical considerations. Unfortunately, such a maximal Ctx is too inefficient in practice as one has to consider the syntactic deducibility problem $\phi \vdash^? s$ for every $t \rightarrow_{\mathcal{R}}^* s$. This is why we rather use the following lighter implementation:

- for every index $0 \leq i \leq N$, and every rule $l \rightarrow r$ in $\mathcal{R}_{i+1} - \mathcal{R}_i$, if $l = D[l_1, \dots, l_n, y_1, \dots, y_{p+q}]$ is a (n, p, q) -decomposition satisfying condition (ii) above for some (arbitrarily chosen) associated context C , then, for every ϕ and σ such that $\phi \vdash l\sigma$, we let

$$\text{Ctx}(\phi \vdash_{\mathcal{R}}^? r\sigma, \Psi, (l, r, D, \sigma)) = C[M_1, \dots, M_{n+p+q}]$$

- where the M_k are fixed recipes such that $(M_i \triangleright l_i\sigma) \in \phi$ for $1 \leq i \leq n$ and $(M_{n+j} \triangleright y_j\sigma) \in \phi$ for $1 \leq j \leq p+q$;
- otherwise, if $\phi \vdash t \downarrow_{\mathcal{R}}$, we let $\text{Ctx}(\phi \vdash_{\mathcal{R}}^? t, \Psi, \alpha)$ be some fixed M such that $M \triangleright_{\phi} t \downarrow_{\mathcal{R}}$;
- in any other case, we let $\text{Ctx}(\phi \vdash_{\mathcal{R}}^? t, \Psi, \alpha) = \perp$.

Using similar ideas as for the proof of Proposition 1, we can show that, for any convergent rewrite system \mathcal{R} , this choice of Ctx is compatible with property (b) of Subsection 4.2, and more generally with completeness, as long as, during saturation, the transformation rules **A** involve the rewrite rules of \mathcal{R}_i with greater priority than those of \mathcal{R}_j , $i < j$. Moreover, when \mathcal{R} is additionally layered, this definition ensures that the procedure never fails. Indeed, using the notations of Figure 1, $\text{Ctx}(\Phi \cup \{z_1 \triangleright z_1, \dots, z_q \triangleright z_q\} \vdash_{\mathcal{R}}^? r\sigma, \Psi, (l, r, D, \sigma)) = \perp$ implies that (ii) is false on D , thus (i) $\text{var}(r) \subseteq \text{var}(l_1, \dots, l_n)$ holds and $(r\sigma) \downarrow_{\mathcal{R}}$ is ground.

Example 6. Any convergent subterm rewrite system \mathcal{R} is layered convergent. Indeed, let $N = 0$ and $\mathcal{R}_1 = \mathcal{R}$. For any $l \rightarrow r$ in \mathcal{R} and for every decomposition $l = D[l_1, \dots, l_n, y_1, \dots, y_p, z_1, \dots, z_q]$, the term r is a subterm of l , thus either $r = C[l_1, \dots, l_n, y_1, \dots, y_p, z_1, \dots, z_q]$ for some context C , or r is a subterm of some l_i thus $\text{var}(r) \subseteq \text{var}(l_1, \dots, l_n)$.

Example 7. Other examples are provided by the theory of homomorphism E_{hom} defined in Section 2.3 as well as the convergent theories of blind signatures E_{blind} and prefix encryption E_{pref} defined by the following sets of equations.

$$\mathcal{E}_{\text{blind}} = \mathcal{E}_{\text{enc}} \cup \left\{ \begin{array}{l} \text{unblind}(\text{blind}(x, y), y) = x \\ \text{unblind}(\text{sign}(\text{blind}(x, y), z), y) = \text{sign}(x, z) \end{array} \right\}$$

$$\mathcal{E}_{\text{pref}} = \mathcal{E}_{\text{enc}} \cup \{ \text{pref}(\text{enc}(\langle x, y \rangle, z)) = \text{enc}(x, z) \}$$

The theory E_{blind} models primitives used in e-voting protocols [17]. The prefix theory represents the property of many chained modes of encryption (e.g. CBC) where an attacker can retrieve any encrypted prefix out of a ciphertext.

Let us check for instance that the prefix theory E_{pref} is layered. Let $N = 1$, \mathcal{R}_1 be the rewrite system obtained from \mathcal{E}_{enc} by orienting the equations from left to right, and $\mathcal{R}_2 = \mathcal{R}_1 \cup \{ \text{pref}(\text{enc}(\langle x, y \rangle, z)) \rightarrow \text{enc}(x, z) \}$. The rewrite rules of \mathcal{R}_1 satisfy the assumptions since \mathcal{R}_1 forms a convergent subterm rewrite system. The additional rule $\text{pref}(\text{enc}(\langle x, y \rangle, z)) \rightarrow \text{enc}(x, z)$ admits three decompositions up to permutation of parameters:

- $l = \text{pref}(l_1)$, in which case $\text{var}(r) \subseteq \text{var}(l_1)$;
- $l = \text{pref}(\text{enc}(l_1, z))$, in which case $\text{enc}(\pi_1(l_1), z) \rightarrow_{\mathcal{R}_1} r$;
- $l = \text{pref}(\text{enc}(\langle x, y \rangle, z))$, in which case $r = \text{enc}(x, z)$.

Verifying that the convergent theories E_{hom} and E_{blind} are layered is similar.

6.2 Termination

In the previous subsection, we described a sufficient criterion for non-failure. To obtain decidability for a given layered convergent theory, there remains only to provide a termination argument. Such an argument is generally easy to develop by hand as we illustrate on the example of the prefix theory. For the case of existing decidability results from [2], such as the theories of blind signature and homomorphic encryption, we also provide a semantic criterion that allows us to directly conclude termination of the procedure.

Proving termination by hand. To begin with, we note that **B** rules always terminate after a polynomial number of steps. Let us write $\xrightarrow{\bullet}^n$ for the relation made of exactly n strict applications of rules ($S \xrightarrow{\bullet} S'$ iff $S \Longrightarrow S'$ and $S \neq S'$).

Proposition 2. *For every states $S = (\Phi, \Psi)$ and S' such that $S \xrightarrow{\bullet}^n S'$ using only **B** rules, n is polynomially bounded in the size of $\text{im}(\Phi)$.*

This is due to the fact that frames are one-to-one and that the rule **B.2** only adds deduction facts $M \triangleright t$ such that t is a subterm of an existing term in Φ . Hence, for proving termination, we observe that it is sufficient to provide a function s mapping each frame Φ to a finite set of terms $s(\Phi)$ including the subterms of $\text{im}(\Phi)$ and such that rule **A.2** only adds deduction facts $M \triangleright t$ satisfying $t \in s(\Phi)$.

For subterm theories, we obtain polynomial termination by choosing $s(\Phi)$ to be the subterms of $\text{im}(\Phi)$ together with the ground right-hand sides of \mathcal{R} .

Proposition 3. *Let \mathbf{E} be a convergent subterm theory. For every $S = (\Phi, \Psi)$ and S' such that $S \xrightarrow{\mathbf{E}}^n S'$, n is polynomially bounded in the size of $\text{im}(\Phi)$.*

To conclude that deduction and static equivalence are decidable in polynomial time [2], we need to show that the deduction facts and the equations are of polynomial size. This requires a DAG representation for terms and visible equations. For our implementation, we have chosen not to use DAGs for the sake of simplicity (and perhaps efficiency) since DAGs require much heavier data structures. However, similar techniques as those described in [2] would apply to implement our procedure using DAGs.

For proving termination of the prefix theory, we let $s(\Phi)$ be the minimal set containing Φ , closed by subterm and such that $\text{enc}(t_1, k) \in s(\Phi)$ whenever $\text{enc}(\langle t_1, t_2 \rangle, k) \in s(\Phi)$. We then deduce that deduction and static equivalence are decidable for the equational theory \mathbf{E}_{pref} , which is a new decidability result.

A criterion to ensure termination. We now provide a semantic criterion that more generally explains why our procedure succeeds on theories previously known to be decidable [2]. This criterion intuitively states that the set of deducible terms from any initial frame φ should be equivalent to a set of *syntactically* deducible terms. Provided that failures are prevented and assuming a *fair* strategy for rule application, we prove that this criterion is a necessary and sufficient condition for our procedure to terminate.

Definition 5 (fair derivation). *An infinite derivation $(\Phi_0, \Psi_0) \Longrightarrow \dots \Longrightarrow (\Phi_n, \Psi_n) \Longrightarrow \dots$ is fair iff along this derivation,*

- (a) *\mathbf{B} rules are applied with greatest priority, and*
- (b) *whenever a \mathbf{A} rule is applicable for some instance $(l \rightarrow r, D, t_1, \dots, t_n, \dots)$, eventually the same instance of rule is applied during the derivation.*

Fairness implies that any deducible term is eventually syntactically deducible.

Lemma 4. *Let $S_0 = (\Phi_0, \Psi_0) \Longrightarrow \dots \Longrightarrow (\Phi_n, \Psi_n) \Longrightarrow \dots$ be an infinite fair derivation from a state S_0 . For every ground term t such that $\Phi_0 \vdash_{\mathbf{E}} t$, either $(\Phi_0, \Psi_0) \Longrightarrow^* \perp$ or there exists i such that $\Phi_i \vdash t \downarrow_{\mathcal{R}}$.*

Proposition 4 (criterion for saturation). *Let φ be an initial frame such that $\text{Init}(\varphi) \not\Rightarrow^* \perp$. The following conditions are equivalent:*

- (i) *There exists a saturated couple (Φ, Ψ) such that $\text{Init}(\varphi) \Longrightarrow^* (\Phi, \Psi)$.*
- (ii) *There exists a (finite) initial frame φ_s such that for every term t , t is deducible from φ modulo \mathbf{E} iff $t \downarrow_{\mathcal{R}}$ is syntactically deducible from φ_s .*
- (iii) *There exists no fair infinite derivation starting from $\text{Init}(\varphi)$.*

Together with the syntactic criterion described in Section 6.1, this criterion (Property (ii)) allows us to prove decidability of deduction and static equivalence for layered convergent theories that belong to the class of *locally stable*

theories defined in [2]. As a consequence, our procedure always saturates for the theories of blind signatures and homomorphic encryption since those theories are layered and have been proved locally stable [2]. Other examples of layered convergent theories enjoying this criterion can be found in [2] (e.g. a theory of addition).

7 Implementation: The YAPA Tool

YAPA is an Ocaml implementation¹ of the saturation procedure presented in Section 4, using by default the optimized function `Ctx` defined in Section 6, and a fair strategy of rule application (see Definition 5).

The tool takes as input an equational theory described by a finite convergent rewrite system, as well as frame definitions and queries. A few optimizations may be activated for subterm theories, e.g. to accelerate normalization. The procedure starts by computing the decompositions of the rewrite system. Provided that the rewrite rules are given in an order compatible with the sets $\mathcal{R}_0 \subseteq \dots \subseteq \mathcal{R}_{N+1}$ of Definition 4, it is able to recognize (fully or partially) layered theories and to pre-compute the associated contexts C related to condition (ii) of this definition, and exploited by the function `Ctx` in use for eliminating failure cases.

We have conducted several experiments on a PC Intel Core 2 Duo at 2.4 GHz with 2 Go RAM for various equational theories (see below) and found that YAPA provides an efficient way to check static equivalence and deducibility.

Equational theory	E_{enc} $n = 10$	E_{enc} $n = 14$	E_{enc} $n = 16$	E_{enc} $n = 18$	E_{enc} $n = 20$	E_{blind}	E_{pref}	E_{hom}	E_{add}
Execution time	< 1s	1,7s	8s	30s	< 3min	< 1s	< 1s	< 1s	< 1s

For the case of E_{enc} , we have run YAPA on the frames $\varphi_n = \{w_1 \triangleright t_n^0, w_2 \triangleright c_0, w_3 \triangleright c_1\}$ and $\varphi'_n = \{w_1 \triangleright t_n^1, w_2 \triangleright c_0, w_3 \triangleright c_1\}$, where $t_0^i = c_i$ and $t_{n+1}^i = \langle \text{enc}(t_n^i, k_n^i), k_n^i \rangle$, $i \in \{0, 1\}$. These examples allow us to increase the (tree, non-DAG) size of the distinguishing tests exponentially, while the sizes of the frames grow linearly. Despite the size of the output, we have observed satisfactory performances for the tool. We have also experimented YAPA on several convergent theories, e.g. E_{blind} , E_{hom} , E_{pref} and the theory of addition E_{add} defined in [2].

In comparison with the tool ProVerif [9,10], here instrumented to check static equivalences, our test samples suggest a running time between one and two orders of magnitude faster for YAPA. Also we did not succeed in making ProVerif terminate on the two theories E_{hom} and E_{add} . Of course, these results are not entirely surprising given that ProVerif is tailored for the more general (and difficult) problem of protocol (in)security under active adversaries. In particular ProVerif's initial preprocessing of the rewrite system appears more substantial than ours and does not terminate on the theories E_{hom} and E_{add} (although termination is guaranteed for linear or subterm-convergent theories [10]).

¹ Freely available at <http://www.lsv.ens-cachan.fr/~baudet/yapa/>

Altogether, these results suggest that YAPA significantly improves the state of the art for checking deducibility and static equivalence under convergent theories, both from practical and theoretical perspectives.

References

1. Abadi, M., Baudet, M., Warinschi, B.: Guessing attacks and the computational soundness of static equivalence. In: Aceto, L., Ingólfssdóttir, A. (eds.) FOSSACS 2006. LNCS, vol. 3921, pp. 398–412. Springer, Heidelberg (2006)
2. Abadi, M., Cortier, V.: Deciding knowledge in security protocols under equational theories. *Theoretical Computer Science* 387(1-2), 2–32 (2006)
3. Abadi, M., Fournet, C.: Mobile values, new names, and secure communication. In: 28th ACM Symposium on Principles of Programming Languages (POPL 2001), pp. 104–115. ACM Press, New York (2001)
4. Anantharaman, S., Narendran, P., Rusinowitch, M.: Intruders with caps. In: Baader, F. (ed.) RTA 2007. LNCS, vol. 4533, pp. 20–35. Springer, Heidelberg (2007)
5. Baudet, M.: Deciding security of protocols against off-line guessing attacks. In: 12th ACM Conference on Computer and Communications Security (CCS 2005), pp. 16–25. ACM Press, New York (2005)
6. Baudet, M.: Sécurité des protocoles cryptographiques : aspects logiques et calculatoires. Thèse de doctorat, LSV, ENS Cachan, France (2007)
7. Baudet, M., Cortier, V., Delaune, S.: YAPA: A generic tool for computing intruder knowledge. Research Report LSV-09-03, Laboratoire Spécification et Vérification, ENS Cachan, France, 28 pages (February 2009)
8. Baudet, M., Cortier, V., Kremer, S.: Computationally sound implementations of equational theories against passive adversaries. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 652–663. Springer, Heidelberg (2005)
9. Blanchet, B.: An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In: 14th Computer Security Foundations Workshop (CSFW 2001), pp. 82–96. IEEE Comp. Soc. Press, Los Alamitos (2001)
10. Blanchet, B., Abadi, M., Fournet, C.: Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming* 75(1), 3–51 (2008)
11. Chevalier, Y., Küsters, R., Rusinowitch, M., Turuani, M.: An NP decision procedure for protocol insecurity with XOR. In: 18th IEEE Symposium on Logic in Computer Science (LICS 2003). IEEE Comp. Soc. Press, Los Alamitos (2003)
12. Comon-Lundh, H., Shmatikov, V.: Intruder deductions, constraint solving and insecurity decision in presence of exclusive or. In: 18th IEEE Symposium on Logic in Computer Science (LICS 2003). IEEE Computer Society Press, Los Alamitos (2003)
13. Corin, R., Doumen, J., Etalle, S.: Analysing password protocol security against off-line dictionary attacks. In: 2nd International Workshop on Security Issues with Petri Nets and other Computational Models (WISP 2004). ENTCS (2004)
14. Cortier, V., Delaune, S.: Deciding knowledge in security protocols for monoidal equational theories. In: Dershowitz, N., Voronkov, A. (eds.) LPAR 2007. LNCS (LNAI), vol. 4790, pp. 196–210. Springer, Heidelberg (2007)
15. Cortier, V., Delaune, S., Lafourcade, P.: A survey of algebraic properties used in cryptographic protocols. *Journal of Computer Security* 14(1), 1–43 (2006)

16. Delaune, S., Jacquemard, F.: A decision procedure for the verification of security protocols with explicit destructors. In: 11th ACM Conference on Computer and Communications Security (CCS 2004), pp. 278–287 (2004)
17. Delaune, S., Kremer, S., Ryan, M.D.: Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security* (to appear) (2008)
18. Lowe, G.: Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In: Margaria, T., Steffen, B. (eds.) TACAS 1996. LNCS, vol. 1055, pp. 147–166. Springer, Heidelberg (1996)
19. Millen, J., Shmatikov, V.: Constraint solving for bounded-process cryptographic protocol analysis. In: 8th ACM Conference on Computer and Communications Security (CCS 2001) (2001)