Computationally sound implementations of equational theories against passive adversaries[☆]

Mathieu Baudet[a], Véronique Cortier[b], Steve Kremer[c]

[a] *DCSSI, France*
[b] *Loria/CNRS & INRIA Lorraine projet Cassis, France*
[c] *LSV/ CNRS & INRIA Saclay projet SECSI & ENS Cachan, France*

Abstract

In this paper we study the link between formal and cryptographic models for security protocols in the presence of passive adversaries. In contrast to other works, we do not consider a fixed set of primitives but aim at results for arbitrary equational theories. We define a framework for comparing a cryptographic implementation and its idealization with respect to various security notions. In particular, we concentrate on the computational soundness of static equivalence, a standard tool in cryptographic pi calculi. We present a soundness criterion, which for many theories is not only sufficient but also necessary. Finally, to illustrate our framework, we establish the soundness of static equivalence for the exclusive OR and a theory of ciphers and lists.

1. Introduction

Today's ubiquity of computer networks increases the need for theoretic foundations for cryptographic protocols. For more than twenty years now, two communities separately developed two families of models. Both views have been very useful in increasing the understanding and quality of security protocol design. On the one hand *formal* or *logical* models have been developed, based on the seminal work of Dolev and Yao [2]. These models view cryptographic operations in a rather abstract and idealized way. On the other hand *cryptographic* or *computational* models [3] are closer to implementations: cryptographic operations are modeled as algorithms manipulating bit-strings. Those models cover a large class of attacks, namely all those implementable by a probabilistic polynomial-time Turing machine.

The advantage of formal models is that security proofs are generally simpler and suitable for automatic procedures, even for complex protocols. Unfortunately, the high degree of abstraction and the limited adversary power raise questions regarding the security offered by such proofs. Potentially, justifying

[☆] An extended abstract of this work was published in the proceedings of the ICALP'05 conference [1].

symbolic proofs with respect to standard computational models has tremendous benefits: protocols can be analyzed using automated tools and still benefit from the security guarantees of the computational model.

For the past few years, a significant research effort has been directed at linking these two approaches. In their seminal work [4], Abadi and Rogaway prove the computational soundness of formal (symmetric) encryption in the case a passive attacker. Since then, many results have been obtained. Each of these results considers a fixed set of primitives, for instance symmetric or public-key encryption. In this paper, we aim at presenting general results for arbitrary equational theories, such as encryption, but also less studied ones, such as groups or exclusive OR. The interest of our approach is not only to develop a general and unified framework for the treatment of cryptographic primitives. Conceptually, it also offers a better understanding of the use of equational theories when modeling the algebraic properties of the primitives. Indeed, for several years, formal models have considered equational theories like the theory of exclusive OR, abelian groups or homomorphic encryption (for a survey on algebraic properties see for instance [5, 6]) in order to model some cryptographic aspects. But it is *a priori* unclear whether "enough" equations have been considered to provide realistic security guarantees. A real attacker might still exploit additional properties of a cryptographic primitive that have not been modeled. Here, we propose a setting and some proof techniques that allow us to formally define and prove that "enough" equations have been considered.

We concentrate on *static equivalence*, a now standard notion originating from the applied pi calculus [7]. Intuitively, static equivalence asks whether an attacker can distinguish between two tuples of messages—later called *frames*—by exhibiting a relation which holds on one tuple but not on the other. Static equivalence provides an elegant means to express security properties on pieces of data, for instance those observed by a passive attacker during the run of a protocol. In the context of active attackers, static equivalence has also been used to characterize process equivalences [7] and off-line guessing attacks [8, 9]. There now exist exact [10] and approximate [11] algorithms to decide static equivalence for a large family of equational theories.

Our first contribution is a general framework for comparing formal and computational models in the presence of a passive attacker. We define the notions of *soundness* and *faithfulness* of a cryptographic implementation with respect to equality, static equivalence and (non-)deducibility. Soundness holds when a formal notion of security has a computational interpretation. For instance, statically equivalent tuples of messages (frames) should be computationally indistinguishable. Conversely, faithfulness holds when every formal attack on a given notion of security can be mapped to an efficient computational attacker. As an illustration, we consider an equational theory modeling Abelian groups with exponents taken over a commutative ring. We show that the soundness of static equivalence implies the hardness of several classical problems in cryptography, notably the decisional Diffie-Hellmann and the RSA problem. Although not completely surprising, this results illustrate well the expressive power of static equivalence defined over tailored equational theories.

Our second contribution is a sufficient criterion for soundness with respect to static equivalence: intuitively the usual computational semantics of terms has to be indistinguishable to an idealized one. We also define and study a useful class of frames, called transparent frames, for arbitrary equational theories. Informally, a frame is transparent if every secret in use is deducible from the frame itself. Transparent frames enjoy notable properties such as a simple characterization of static equivalence and—in the case of uniform distributions—the fact that two statically equivalent transparent frames always yield the same concrete distribution, that is, are indistinguishable in the sense of information theory. This study of transparent frames allows us to exhibit a class of equational theories for which our soundness criterion is necessary.

Our third contribution consists in applying our framework to obtain two first soundness results for static equivalence. The first equational theory that we consider deals with the exclusive OR. This simple but important primitive has been largely used in cryptographic constructions such as the One-Time Pad and in protocols (see [6] for examples). Interestingly, our proof of soundness reflects the unconditional security (in the information-theoretic sense) of the One-Time Pad [12]. Second we consider a theory of symmetric encryption and lists. The result is similar in spirit to the one of Abadi and Rogaway [4]. However, we consider deterministic, length-preserving, symmetric encryption schemes—also known as pseudo-random permutations or ciphers, while Abadi and Rogaway consider probabilistic, symmetric encryption. This choice is motivated by famous examples of ciphers such as DES or AES. In both examples, the specificity of our work is to prove the soundness of a standard formal notion, static equivalence, rather than that of a specialized relation.

Related work.. The study of the link between the formal and the computational approaches for cryptographic protocols started with the seminal work of Abadi and Rogaway [4], in a passive setting. There have been many extensions to the work of Abadi and Rogaway in the passive case, such as studying completeness [13], considering deterministic encryption [14] (a more detailed comparison is provided below), One-Time pad, length-revealing and same-key revealing encryption [12] or allowing composed keys [15] and key-cycles [16].

The first results in an active setting were achieved by Backes, Pfitzmann, and Waidner [17, 18, 19]. These works prove the soundness of a rich language including digital signatures, public-key and symmetric key encryption in the presence of an active attacker for several kind of security properties. Quite similar results were established in more abstract and classical Dolev-Yao models for asymmetric encryption and signatures [20, 21]. While more easily amendable to full automation, these results do not offer universal composability guarantees like the previous ones. However, Canetti and Herzog [22] have recently obtained a similar soundness theorem for a restricted class of protocols—mutual authentication and key exchange protocols using only public-key encryption—which does offer strong composability properties in the universal composability framework. Laud [23] presents an automated procedure for computationally sound proofs of confidentiality in the case of an active attacker and symmetric

encryption when the number of sessions is bounded. Datta et al. [24] introduce a symbolic logic that allows cryptographically sound security proof. Recently, Blanchet [25] proposed a computationally sound mechanized prover that relies directly on games transformations, a proof technique commonly used in the cryptographic setting.

Except [25], the previously mentioned results are all dedicated to some fixed set of cryptographic primitives. Here, our goal is not restricted to obtaining some particular soundness result for a given set of primitives and security properties. Rather, we aim at developing a general setting to reason about the adequacy of abstract functional symbols equipped with an equational theory and their corresponding cryptographic implementations. To the best of our knowledge, this approach is new and distinct from existing work. We now discuss some related work concerning the two theories (exclusive OR as well as ciphers and lists) that we have considered to illustrate our framework.

Regarding the soundness of exclusive OR, Backes and Pfitzmann [26] have independently shown an impossibility result in the framework of reactive simulatability, in the presence of an active adversary. They also present a soundness result in the presence of a passive adversary. While we consider the application of exclusive OR only to pure random values, Backes and Pfitzmann deal with arbitrary payloads. It is however not clear how the framework of reactive simulatability in the presence of a passive adversary compares to our framework based on static equivalence.

Concerning the theory of ciphers and list, Laud [14] presents soundness results in the style of Abadi and Rogaway for ciphers. While these results are close to ours, Laud's notion of formal equivalence is apparently more pessimistic than ours regarding the secrecy of encryption keys. For instance, as opposed to [14], we consider that the encryption of a fresh random value by a known key is indistinguishable from a random value—that is, formally, the pair $(\mathsf{enc}(n,k),k)$ is indistinguishable from (n',k). The reason is that, in the absence of tags, each encryption key of a cipher yields a permutation on the space of values. Therefore, if n follows the uniform distribution, such as in our implementation (Section 5.2), so does the term $\mathsf{enc}(n,k)$. Provided a suitable set of equations, static equivalence naturally accounts for this property, whereas there seems to be no natural and immediate way to express the same equivalences using patterns in the style of Abadi and Rogaway. In some sense, the work of Abadi and Warinschi [27] can be seen as an attempt to do so on a fragment of equivalences modeling guessing attacks. Recently, the techniques developed in the present paper have been applied successfully by Abadi, Baudet, and Warinschi [28] to generalize the ideas of [27] and justify a modeling of guessing attacks purely based on static equivalence.

In [14], Laud provides a computationally sound proof system handling both ciphers and exclusive OR in the presence of a passive attacker. This proof system is used to prove the security of several encryption modes including CBC. This approach differs from the one developed here as it aims at direct cryptographic proofs of security (much as in [23, 25]). In comparison, our approach (as in [4, 12, 15, 16, 17, 18, 19, 13, 20, 21]) aims to exhibit a class of protocols for which

the absence of formal attacks entails the existence of a computational proof of security.

Further related work.. Since the publication of a preliminary version [1] of this article, several papers have addressed the computational soundness of static equivalence. As already mentioned, Abadi, Baudet, and Warinschi [28] study resistance against offline guessing attacks modelled in terms of static equivalence and use the framework developed in this paper to show the soundness of an equational theory including ciphers, symmetric and asymmetric encryption. In [29], Bana, Mohassel and Stegers argue that the notion of static equivalence is too coarse and not sound for many interesting equational theories. They introduce a general notion of formal indistinguishability relation. This highlights that soundness of static equivalence only holds for a restricted set of well-formed frames (in the same vein Abadi and Rogaway used restrictions to forbid key cycles). They illustrate the unsoundness of static equivalence for modular exponentiation. More recently, Kremer and Mazaré [30] use our framework to define soundness of static equivalence in the presence of an adaptive, rather than purely passive, adversary. They show soundness results of static equivalence for an equational theory modelling modular exponentiation (for a class of well-formed frames, hence not contradicting [29]), as well as symmetric encryption with composed keys which can be computed using modular exponentiation or exclusive or.

The active version of static equivalence is the observational equivalence relation introduced by Milner and Hoare in the early 80s. Intuitively, two processes are equivalent if an observer cannot tell the difference between the two processes. The observer can in particular intercept and send messages to the processes. Comon-Lundh and Cortier [31] have recently shown that observational equivalence between processes in a fragment of the applied pi-calculus [32] implies cryptographic indistinguishability against active attackers, in the context of symmetric encryption. They use an extended version of soundness of static equivalence (called tree soundness) as a key step in their proof.

Outline of the paper.. In the next section, we introduce our abstract and concrete models together with the notions of indistinguishability. We then define the notions of soundness and faithfulness and illustrate some consequences of soundness with respect to static equivalence on groups. In Section 4, we define the ideal semantics of abstract terms, present our soundness criterion, and prove it necessary for a large family of equational theories. As an illustration (Section 5), we prove the soundness for the theories modeling exclusive OR, as well as ciphers and lists. We then conclude in Section 6. An appendix contains detailed proofs of formal lemmas related to static equivalence.

2. Modeling cryptographic primitives with abstract algebras

In this section we introduce some notations and set our abstract and concrete models.

2.1. Abstract algebras

Our abstract models—called *abstract algebras*—consist of term algebras defined over a many-sorted first-order signature and equipped with equational theories.

Specifically, a *signature* $(\mathcal{S}, \mathcal{F})$ is made of a set of *sorts* $\mathcal{S}$, with elements denoted by s, $s_1 \ldots$, and a set $\mathcal{F}$ of *symbols*, written f, $f_1 \ldots$, together with arities of the form $\mathrm{ar}(f) = s_1 \times \ldots \times s_k \to s$ $(k \geq 0)$. Symbols that take $k = 0$ arguments are called *constants*; their arity is simply written s. We fix a set $\mathcal{N}$ of *names*, written a, $b \ldots$, and a set $\mathcal{X}$ of *variables* x, $y \ldots$ We assume that names and variables are given with sorts, and that an infinite number of names and variables are available for each sort. The set of *terms of sort* s is defined inductively by

T	::=		term of sort s
	$\mid$	x	variable x of sort s
	$\mid$	a	name a of sort s
	$\mid$	$f(T_1, \ldots, T_k)$	application of symbol $f \in \mathcal{F}$

where for the last case, we further require that T_i is a term of some sort s_i and $\mathrm{ar}(f) = s_1 \times \ldots \times s_k \to s$. We write $\mathrm{var}(T)$ and $\mathrm{names}(T)$ for the set of variables and names occurring in T, respectively. A term T is *ground* or *closed* iff $\mathrm{var}(T) = \emptyset$. We may write $\mathrm{var}(T_1, \ldots, T_k)$ instead of $\mathrm{var}(\{T_1, \ldots, T_k\})$ and similarly for names.

A *context* C is a term with holes, or (more formally) a term with distinguished variables. When C is a context with n distinguished variables x_1, ..., x_n, we may write $C[x_1, \ldots, x_n]$ instead of C in order to show the variables, and when T_1, ..., T_n are terms we may also write $C[T_1, \ldots, T_n]$ for the result of replacing each variable x_i with the corresponding term T_i.

Substitutions are written $\sigma = \{x_1 \mapsto T_1, \ldots, x_n \mapsto T_n\}$ with domain $\mathrm{dom}(\sigma) = \{x_1, \ldots, x_n\}$. We only consider *well-sorted* substitutions, that is, substitutions $\sigma = \{x_1 \mapsto T_1, \ldots, x_n \mapsto T_n\}$ for which x_i and T_i have the same sort. Such a σ is *closed* iff all of the T_i are closed. We let $\mathrm{var}(\sigma) = \bigcup_i \mathrm{var}(T_i)$, $\mathrm{names}(\sigma) = \bigcup_i \mathrm{names}(T_i)$, and extend the notations var(.) and names(.) to tuples and sets of terms and substitutions in the obvious way. The application of a substitution σ to a term T is written $\sigma(T) = T\sigma$. If p is a position of T, the expression $T|_p$ denotes the subterm of T at the position p. The expression $T[T']_p$ denotes the term obtained after replacing the subterm in position p of T with T'.

Symbols in $\mathcal{F}$ are intended to model cryptographic primitives, whereas names in $\mathcal{N}$ are used to model secrets, that is, concretely random numbers. The intended behavior of the primitives is described by an equational theory E, that is, an equivalence relation on terms (also written $=_E$) compatible with applications of symbols and well-sorted substitutions:

- for every k-ary symbol f and terms $t_1, \ldots, t_k, t'_1, \ldots, t'_k$ of the appropriate sorts, $\forall i,\ t_i =_E t'_i$ implies that $f(t_1, \ldots, f_k) =_E f(t'_1, \ldots, f'_k)$;

- for every well-sorted substitution σ and terms t, t', if $t =_E t'$ then $t\sigma =_E t'\sigma$.

In the sequel we further require that E is stable under (well-sorted) substitution of names. All the equational theories that we consider in this paper satisfy these properties. For instance, symmetric and deterministic encryption is modeled by the theory E_{enc} generated by the classical equation $E_{\mathsf{enc}} = \{\mathsf{dec}(\mathsf{enc}(x, y), y) = x\}$.

A symbol f is *free* with respect to an equational theory E iff there exists a set of equations F generating E such that f does not occur in F. A sort s is *degenerated* in E iff all terms of sort s are equal modulo E.

It is often useful to orient equations and work with *rewriting rules* instead of the equational theory. Formally, a *rewriting rule* is an expression $l \to r$ where l and r are two terms of the same sort. Given a set of rewriting rules $\mathcal{R}$ (called *rewriting system*), we write $T \to_{\mathcal{R}} T'$ if there exists a rule $l \to r \in \mathcal{R}$, a position p and a (well-sorted) substitution σ such that $T|_p = l\sigma$ and $T' = T[r\sigma]_p$. We write $\to^*_{\mathcal{R}}$ for the reflexive and transitive closure of $\to_{\mathcal{R}}$, and $=_{\mathcal{R}}$ for its reflexive, symmetric and transitive closure.

Given an equational theory E and a rewriting system $\mathcal{R}$, we write $\to_{\mathcal{R}/E}$ for the relation $=_E \to_{\mathcal{R}} =_E$. We define $\to^*_{\mathcal{R}/E}$ and $=_{\mathcal{R}/E}$ similarly as above. $\mathcal{R}$ is *E-terminating* iff $\to_{\mathcal{R}/E}$ admits no infinite sequence of reductions $T_0 \to_{\mathcal{R}/E} T_1 \to_{\mathcal{R}/E} \ldots T_n \to_{\mathcal{R}/E} \ldots$. It is *$E$-confluent* iff for every $T \to^*_{\mathcal{R}/E} T_1$ and $T \to^*_{\mathcal{R}/E} T_2$, there exist T'_1 and T'_2 such that $T_1 \to^*_{\mathcal{R}/E} T'_1$, $T_2 \to^*_{\mathcal{R}/E} T'_2$, and $T'_1 =_E T'_2$. Finally, $\mathcal{R}$ is *E-convergent* iff it is both E-terminating and E-confluent. When E is the syntactic equality, this yields the usual notions of termination, confluence and convergence.

2.2. Frames, deducibility and static equivalence

We use frames [7, 10] to represent sequences of messages observed by an attacker, for instance during the execution of a protocol. Formally, a (closed) *frame* is an expression $\varphi = \nu\tilde{a}.\{x_1 = T_1, \ldots, x_n = T_n\}$ where $\tilde{a}$ is a set of *bound (or restricted) names*, and for each i, T_i is a closed term of the same sort as x_i.

For simplicity, we only consider (closed) frames $\varphi = \nu\tilde{a}.\{x_1 = T_1, \ldots, x_n = T_n\}$ which restrict every name in use, that is, for which $\tilde{a} = \mathrm{names}(T_1, \ldots, T_n)$. A name a may still be disclosed explicitly by adding a mapping $x_a = a$ to the frame. Thus we tend to assimilate such frames φ to their *underlying substitutions* $\sigma = \{x_1 \mapsto T_1, \ldots, x_n \mapsto T_n\}$.

Definition 1 (Deducibility). A (closed) term T is *deducible* from a frame φ in an equational theory E, written $\varphi \vdash_E T$, iff there exists a term M such that $\mathrm{var}(M) \subseteq \mathrm{dom}(\varphi)$, $\mathrm{names}(M) \cap \mathrm{names}(\varphi) = \emptyset$, and $M\varphi =_E T$.

In what follows, again for simplicity, we only consider deducibility problems $\varphi \vdash_E T$ such that $\mathrm{names}(T) \subseteq \mathrm{names}(\varphi)$.

Consider for instance the theory E_{enc} and the frame $\varphi_1 = \nu k_1, k_2, k_3, k_4.\{x_1 = \mathsf{enc}(k_1, k_2), x_2 = \mathsf{enc}(k_4, k_3), x_3 = k_3\}$: the name k_4 is deducible from φ_1 since $\mathsf{dec}(x_2, x_3)\varphi_1 =_{E_{\mathsf{enc}}} k_4$ but neither k_1 nor k_2 are deducible.

Deducibility is not always sufficient to account for the knowledge of an attacker. For instance, it lacks partial information on secrets. Indeed, if we consider a naive vote protocol where agents simply send their vote (0 or 1) encrypted under some key, the security problem is not whether an attacker can learn the values of 0 or 1, but rather whether an attacker can tell the difference between a message that contains the vote 0 and a message that contains the vote 1. That is why another classical notion in formal methods is *static equivalence*.

Definition 2 (Static equivalence). Two frames φ_1 and φ_2 are *statically equivalent* in a theory E, written $\varphi_1 \approx_E \varphi_2$, iff $\mathrm{dom}(\varphi_1) = \mathrm{dom}(\varphi_2)$, and for all terms M and N such that $\mathrm{var}(M, N) \subseteq \mathrm{dom}(\varphi_1)$ and $\mathrm{names}(M, N) \cap \mathrm{names}(\varphi_1, \varphi_2) = \emptyset$, $M\varphi_1 =_E N\varphi_1$ if and only if $M\varphi_2 =_E N\varphi_2$.

For instance, the two frames $\nu k.\{x = \mathsf{enc}(0, k)\}$ and $\nu k.\{x = \mathsf{enc}(1, k)\}$ are statically equivalent with respect to E_{enc}. However the two frames

$$\nu k.\{x = \mathsf{enc}(0, k),\ y = k\} \text{ and } \nu k.\{x = \mathsf{enc}(1, k),\ y = k\}$$

are not (consider the test $\mathsf{dec}(x, y) \stackrel{?}{=} 0$), although the set of terms that can be deduced from both frames is the same (0 and 1 are two constants known by the attacker).

2.3. Concrete semantics

We now give terms and frames a concrete semantics, parameterized by an implementation of the primitives. Provided a set of sorts $\mathcal{S}$ and a set of symbols $\mathcal{F}$ as above, a *$(\mathcal{S}, \mathcal{F})$-computational algebra* A consists of

- a non-empty set of bit-strings $[\![s]\!]_A \subseteq \{0, 1\}^*$ for each sort $s \in \mathcal{S}$;
- an effective procedure implementing a function $[\![f]\!]_A : [\![s_1]\!]_A \times \ldots \times [\![s_k]\!]_A \to [\![s]\!]_A$ for each symbol $f \in \mathcal{F}$ with $\mathrm{ar}(f) = s_1 \times \ldots \times s_k \to s$;
- an effective procedure for deciding a congruence $=_{A,s}$ for each sort s, in order to check the equality of elements in $[\![s]\!]_A$ (the same element may be represented by different bit-strings); by congruence, we mean a reflexive, symmetric, transitive relation such that $e_1 =_{A,s_1} e'_1, \ldots, e_k =_{A,s_k} e'_k \Rightarrow [\![f]\!]_A(e_1, \ldots, e_k) =_{A,s} [\![f]\!]_A(e'_1, \ldots, e'_k)$ (in the remaining we often omit s and write $=_A$ for $=_{A,s}$);
- an effective procedure to draw random elements from $[\![s]\!]_A$; we denote such a drawing by $x \stackrel{R}{\leftarrow} [\![s]\!]_A$; the drawing may not follow a uniform distribution, but no $=_{A,s}$-equivalence class should have probability 0.

Assume a fixed $(\mathcal{S}, \mathcal{F})$-computational algebra A. We associate to each (closed) frame $\varphi = \{x_1 = T_1, \ldots, x_n = T_n\}$ a distribution $\psi = [\![\varphi]\!]_A$, of which the drawings $\widehat{\psi} \stackrel{R}{\leftarrow} \psi$ are computed as follows:

1. for each name a of sort s appearing in $T_1, \ldots, T_n$, draw a value $\widehat{a} \stackrel{R}{\leftarrow} [\![s]\!]_A$;

2. for each x_i $(1 \leq i \leq n)$ of sort s_i, compute $\widehat{T_i} \in [\![s_i]\!]_A$ recursively on the structure of terms: $\widehat{f(T'_1, \ldots, T'_m)} = [\![f]\!]_A(\widehat{T'_1}, \ldots, \widehat{T'_m})$; using the values $\widehat{a}$ defined at step 1 for names.
3. return the value $\widehat{\psi} = \{x_1 = \widehat{T_1}, \ldots, x_n = \widehat{T_n}\}$.

Such values $\phi = \{x_1 = e_1, \ldots, x_n = e_n\}$ with $e_i \in [\![s_i]\!]_A$ are called *concrete frames*. We extend the notation $[\![.]\!]_A$ to tuples of closed terms in the natural way: $e_1, \ldots, e_n \xleftarrow{R} [\![T_1, \ldots, T_n]\!]_A$ denotes the drawing

$$\{x_1 = e_1, \ldots, x_n = e_n\} \xleftarrow{R} [\![\{x_1 = T_1, \ldots, x_n = T_n\}]\!]_A$$

for appropriate variables x_1, ..., x_n. We also generalize the notation to (tuples of) terms with variables, by specifying a concrete value for each of them: $[\![\cdot]\!]_{A,\{x_1=e_1,\ldots,x_n=e_n\}}$. Notice that when a term or a frame contains no names, the translation is deterministic; in this case, we use the same notation to denote the distribution and its unique value.

In the rest of the paper we focus on asymptotic notions of cryptographic security and consider families of computational algebra (A_η) indexed by a complexity parameter $\eta \geq 0$. (This parameter η might be thought as the size of keys and other secret values.) The *concrete semantics* of a frame φ is a family of distributions over concrete frames $([\![\varphi]\!]_{A_\eta})$. We only consider families of computational algebras (A_η) such that the algebraic operations (i.e. the functions associated to symbols, the congruence relation $=_{A,s}$, and the drawing functions) are computable by uniform, probabilistic polynomial-time algorithms in the complexity parameter η. This ensures that the concrete semantics of every (fixed) term or frame is efficiently computable (in the same sense).

Families of distributions (*ensembles*) over concrete frames benefit from the usual notion of cryptographic indistinguishability. Intuitively, two families of distributions (ψ_η) and (ψ'_η) are *indistinguishable*, written $(\psi_\eta) \approx (\psi'_\eta)$, iff no probabilistic polynomial-time adversary $\mathcal{A}$ can guess whether he is given a sample from ψ_η or ψ'_η with a probability significantly greater than $\frac{1}{2}$. Formally, we ask the *advantage* of $\mathcal{A}$,

$$\mathrm{Adv}^{\mathrm{IND}}(\mathcal{A}, \eta, \psi_\eta, \psi'_\eta) = \mathbb{P}[\widehat{\psi} \xleftarrow{R} \psi_\eta \ : \ \mathcal{A}(\eta, \widehat{\psi}) = 1] - \mathbb{P}[\widehat{\psi} \xleftarrow{R} \psi'_\eta \ : \ \mathcal{A}(\eta, \widehat{\psi}) = 1]$$

to be a *negligible* function of η. We recall that a function f is said *negligible* if for any integer $n > 0$, there exists η_0 such that $f(\eta) \leq \eta^{-n}$ for any $\eta \geq \eta_0$. (Note that we regard negative functions as negligible here.)

A function $f(\eta)$ is *overwhelming* iff $1 - f(\eta)$ is negligible. A family of distributions (ψ_η) is *collision-free* (with respect to the family of congruences $=_{A_\eta}$) iff the probability of collision between two random elements from ψ_η, that is, $\mathbb{P}[\, e_1, e_2 \xleftarrow{R} \psi_\eta \ : \ e_1 =_{A_\eta} e_2]$, is a negligible function of η. Note that, by classical properties of probability, this is equivalent to requiring that the probability of sampling any given e_0 from ψ_η (modulo $=_{A_\eta}$) is negligible, that is, the function $\sup_{e_0} \mathbb{P}\left[e \xleftarrow{R} \psi_\eta \ : \ e =_{A_\eta} e_0\right]$ is bounded by a negligible function of η.

By convention, the adversaries considered in this paper are given access implicitly to the complexity parameter η and to as many fresh random coins as needed.

3. Comparing abstract and computational algebras

In the previous section we have defined abstract and computational algebras. We now relate formal notions such as equality, (non-)deducibility and static equivalence to their computational counterparts, that is, equality, one-wayness and indistinguishability.

3.1. Soundness and faithfulness

We introduce the notions of sound and faithful computational algebras with respect to the formal relations studied here: equality, static equivalence and deducibility.

Let E be an equational theory. A family of computational algebras (A_η) is

- $=_E$*-sound* iff for every closed terms T_1, T_2 of the same sort, $T_1 =_E T_2$ implies that $\mathbb{P}[\, e_1, e_2 \xleftarrow{R} [\![T_1, T_2]\!]_{A_\eta} \,:\, e_1 =_{A_\eta} e_2]$ is overwhelming;
- $=_E$*-faithful* iff for every closed terms T_1, T_2 of the same sort, $T_1 \neq_E T_2$ implies that $\mathbb{P}[\, e_1, e_2 \xleftarrow{R} [\![T_1, T_2]\!]_{A_\eta} \,:\, e_1 =_{A_\eta} e_2]$ is negligible;
- $\approx_E$*-sound* iff for every frames φ_1, φ_2 with the same domain, $\varphi_1 \approx_E \varphi_2$ implies that $([\![\varphi_1]\!]_{A_\eta}) \approx ([\![\varphi_2]\!]_{A_\eta})$;
- $\approx_E$*-faithful* iff for every frames φ_1, φ_2 of the same domain, $\varphi_1 \not\approx_E \varphi_2$ implies that there exists a polynomial-time adversary $\mathcal{A}$ for distinguishing concrete frames, such that $\mathrm{Adv}^{\mathrm{IND}}(\mathcal{A}, \eta, [\![\varphi_1]\!]_{A_\eta}, [\![\varphi_2]\!]_{A_\eta})$ is overwhelming;
- $\not\vdash_E$*-sound* iff for every frame φ and closed term T such that $\mathrm{names}(T) \subseteq \mathrm{names}(\varphi)$, $\varphi \not\vdash_E T$ implies that for each polynomial-time adversary $\mathcal{A}$, $\mathbb{P}[\phi, e \xleftarrow{R} [\![\varphi, T]\!]_{A_\eta} : \mathcal{A}(\phi) =_{A_\eta} e]$ is negligible;
- $\not\vdash_E$*-faithful* iff for every frame φ and closed term T such that $\mathrm{names}(T) \subseteq \mathrm{names}(\varphi)$, $\varphi \vdash_E T$ implies that there exists a polynomial-time adversary $\mathcal{A}$ such that $\mathbb{P}[\phi, e \xleftarrow{R} [\![\varphi, T]\!]_{A_\eta} \,:\, \mathcal{A}(\phi) =_{A_\eta} e]$ is overwhelming.

Sometimes, it is possible to prove stronger notions of soundness that hold without restriction on the computational power of adversaries. In particular, (A_η) is

- *unconditionally* $=_E$*-sound* iff for every closed terms T_1, T_2 of the same sort, $T_1 =_E T_2$ implies that $\mathbb{P}[\, e_1, e_2 \xleftarrow{R} [\![T_1, T_2]\!]_{A_\eta} \,:\, e_1 =_{A_\eta} e_2] = 1$;
- *unconditionally* $\approx_E$*-sound* iff for every frames φ_1, φ_2 with the same domain, $\varphi_1 \approx_E \varphi_2$ implies $([\![\varphi_1]\!]_{A_\eta}) = ([\![\varphi_2]\!]_{A_\eta})$;

- *unconditionally $\not\vdash_E$-sound* iff for every frame φ and closed term T such that $\mathrm{names}(T) \subseteq \mathrm{names}(\varphi)$ and $\varphi \not\vdash_E T$, the drawings for φ and T are independent: for all ϕ_0, e_0, $\mathbb{P}[\phi_0, e_0 \xleftarrow{R} [\![\varphi, T]\!]_{A_\eta}] = \mathbb{P}[\phi_0 \xleftarrow{R} [\![\varphi]\!]_{A_\eta}] \times \mathbb{P}[e_0 \xleftarrow{R} [\![T]\!]_{A_\eta}]$, and the drawing $(\xleftarrow{R} [\![T]\!]_{A_\eta})$ is collision-free.

The fact that the first two unconditional notions are stronger than their computational counterparts is clear from the definitions. As for the unconditional $\not\vdash_E$-soundness, observe that if the drawings for φ and T are independent, and the drawing $(\xleftarrow{R} [\![T]\!]_{A_\eta})$ is collision-free, then any adversary $\mathcal{A}$ has negligible probability of retrieving the value of T:

$$\begin{aligned}\mathbb{P}[\phi, e \xleftarrow{R} [\![\varphi, T]\!]_{A_\eta} : \mathcal{A}(\phi) =_{A_\eta} e] \\ &= \mathbb{P}[\phi \xleftarrow{R} [\![\varphi]\!]_{A_\eta}, e \xleftarrow{R} [\![T]\!]_{A_\eta} : \mathcal{A}(\phi) =_{A_\eta} e] \\ &\leq \sup_{e_0} \mathbb{P}[e \xleftarrow{R} [\![T]\!]_{A_\eta} : e =_{A_\eta} e_0]\end{aligned}$$

Generally, (unconditional) $=_E$-soundness is given by construction. Indeed true formal equations correspond to the expected behavior of primitives and should hold in the concrete world with overwhelming probability. The other criteria are however more difficult to fulfill. Therefore it is often interesting to restrict frames to *well-formed* ones in order to achieve soundness or faithfulness: for instance Abadi and Rogaway [4] do forbid encryption cycles (see Section 5.2).

It is worth noting that the notions of soundness and faithfulness introduced above are not independent.

Proposition 1. *Let (A_η) be a $=_E$-sound family of computational algebras. Then*

1. *(A_η) is $\not\vdash_E$-faithful;*
2. *if (A_η) is also $=_E$-faithful, (A_η) is $\approx_E$-faithful.*

Proof.

1. Suppose $\mathrm{names}(T) \subseteq \mathrm{names}(\varphi)$ and $\varphi \vdash_E T$, that is, there exists M such that $\mathrm{var}(M) \subseteq \mathrm{dom}(\varphi)$, $\mathrm{names}(M) \cap \mathrm{names}(\varphi) = \emptyset$, and $M\varphi =_E T$. We define an adversary $\mathcal{A}$ which can deduce $[\![T]\!]$ from $[\![\varphi]\!]$ as follows: given the concrete frame $\phi = \{x_i = e_i\}$, $\mathcal{A}$ returns a sample $e \xleftarrow{R} [\![M]\!]_{A_\eta,\phi}$. As $(A_\eta)_{\eta\geq 0}$ is $=_E$-sound and $\mathrm{names}(T) \subseteq \mathrm{names}(\varphi)$, $\mathcal{A}$'s probability of success is greater than 1 minus a negligible function.

2. Suppose $\varphi_1 \not\approx_E \varphi_2$: there exist two terms M and N such that $\mathrm{var}(M, N) \subseteq \mathrm{dom}(\varphi_1)$, $\mathrm{names}(M, N) \cap \mathrm{names}(\varphi_1, \varphi_2) = \emptyset$, and for instance $M\varphi_1 =_E N\varphi_1$ whereas $M\varphi_2 \neq_E N\varphi_2$. Let $\mathcal{A}$ be the adversary that tests, given η and ϕ, whether $[\![M]\!]_{A_\eta,\phi} =_{A_\eta} [\![N]\!]_{A_\eta,\phi}$, and returns the result of the test. $\mathcal{A}$ runs in polynomial-time and by $=_E$-soundness and $=_E$-faithfulness, its advantage is 1 minus a negligible function. □

For many theories, we have that $\approx_E$-soundness implies all the other notions of soundness and faithfulness. This emphasizes the importance of $\approx_E$-soundness and provides an additional motivation for its study. As an illustration, let us consider an arbitrary theory which includes keyed hash functions.

Proposition 2. *Let (A_η) be a family of $\approx_E$-sound computational algebras. Assume that free binary symbols $\mathsf{h}_s : s \times Key \to Hash$ are available for every sort s, where the sort Key is not degenerated in E, and the drawing of random elements for the sort $Hash$, $(\xleftarrow{R} [\![Hash]\!]_{A_\eta})$, is collision-free. Then*

1. *(A_η) is $=_E$-faithful;*
2. *(A_η) is $\not\vdash_E$-sound;*
3. *Assume the implementations for the symbols h_s are* collision-resistant*, that is, assume that for all T_1, T_2 of sort s, given a fresh name k of sort Key, the quantity*

$$\mathbb{P}\left[e_1, e_2, e_1', e_2' \xleftarrow{R} [\![T_1, T_2, \mathsf{h}_s(T_1, k), \mathsf{h}_s(T_2, k)]\!]_{A_\eta} \ : \ e_1 \neq_{A_\eta} e_2,\ e_1' =_{A_\eta} e_2'\right]$$

is negligible. Then (A_η) is $=_E$-sound, $\not\vdash_E$-faithful and $\approx_E$-faithful.

Proof.

1. Let T_1, T_2 be two terms of sort s such that $T_1 \neq_E T_2$. Consider the frame $\varphi = \{x_1 = \mathsf{h}_s(T_1, k),\ x_2 = \mathsf{h}_s(T_2, k)\}$ where k is a fresh name of sort Key. As $T_1 \neq_E T_2$ and h_s is free, we have $\varphi \approx_E \{x_1 = n,\, x_2 = n'\}$ where n, n' are two distinct fresh names of sort $Hash$ (Proposition 17 of Appendix A). By assumption, this entails $[\![\varphi]\!] \approx [\![\{x_1 = n,\, x_2 = n'\}]\!]$. In particular, since $(\xleftarrow{R} [\![Hash]\!]_{A_\eta})$ is collision-free, the quantity

$$\mathbb{P}\left[e_1, e_2 \xleftarrow{R} [\![T_1,\, T_2]\!]_{A_\eta} \ : \ e_1 =_{A_\eta} e_2\right] \leq \mathbb{P}\left[e_1', e_2' \xleftarrow{R} [\![\mathsf{h}_s(T_1, k),\, \mathsf{h}_s(T_2, k)]\!]_{A_\eta} \ : \ e_1' =_{A_\eta} e_2'\right]$$

   is negligible.

2. Let φ be a frame and T a closed term of sort s such that $\mathrm{names}(T) \subseteq \mathrm{names}(\varphi)$ and $\varphi \not\vdash_E T$. We let $\varphi_0 = \varphi \cup \{x = \mathsf{h}_s(T, k), y = k\}$ and $\varphi_1 = \varphi \cup \{x = n, y = k\}$ where x, y are fresh variables, k is a fresh name of sort Key, n is a fresh name of sort $Hash$. As $\varphi \not\vdash_E T$, we have $\varphi_0 \approx_E \varphi_1$ (Proposition 18 of Appendix A). Thus by assumption, $[\![\varphi_0]\!] \approx [\![\varphi_1]\!]$.

   By contradiction, suppose that there exists a polynomial-time adversary $\mathcal{A}$ able to deduce $[\![T]\!]$ from $[\![\varphi]\!]$ concretely with non-negligible probability of success. We build an adversary $\mathcal{B}$ that distinguishes between $[\![\varphi_0]\!]$ and $[\![\varphi_1]\!]$ as follows: let ϕ be the sample from $[\![\varphi_b]\!]_\eta$ to be analyzed, where $b \in \{0, 1\}$. Let $\widehat{T}$ be the answer of $\mathcal{A}$ when given the restriction of ϕ

to $\mathrm{dom}(\varphi)$. $\mathcal{B}$ returns 0 if $x\phi =_{A_\eta} [\![\mathsf{h}_s]\!]_{A_\eta}(\widehat{T}, y\phi)$, and 1 otherwise. By definition, the advantage of $\mathcal{B}$ is

$$\begin{aligned}
&\mathbb{P}[\phi \xleftarrow{R} [\![\varphi_0]\!]_\eta \ : \ \mathcal{B}(\eta,\phi)=0] - \mathbb{P}[\phi \xleftarrow{R} [\![\varphi_1]\!]_\eta \ : \ \mathcal{B}(\eta,\phi)=0] \\
&= \quad \mathbb{P}[\phi \xleftarrow{R} [\![\varphi_0]\!]_\eta; \widehat{T} \xleftarrow{R} \mathcal{A}(\phi|_{\mathrm{dom}(\varphi)}) \ : \ x\phi =_{A_\eta} [\![\mathsf{h}_s]\!]_{A_\eta}(\widehat{T}, y\phi)] \\
&\quad - \mathbb{P}[\phi \xleftarrow{R} [\![\varphi_1]\!]_\eta; \widehat{T} \xleftarrow{R} \mathcal{A}(\phi|_{\mathrm{dom}(\varphi)}) \ : \ x\phi =_{A_\eta} [\![\mathsf{h}_s]\!]_{A_\eta}(\widehat{T}, y\phi)] \\
&\geq \quad \mathbb{P}[\phi, e \xleftarrow{R} [\![\varphi_0, T]\!]_\eta; \widehat{T} \xleftarrow{R} \mathcal{A}(\phi|_{\mathrm{dom}(\varphi)}) \ : \ \widehat{T} = e)] \\
&\quad - \mathbb{P}[\phi \xleftarrow{R} [\![\varphi_1]\!]_\eta; \widehat{T} \xleftarrow{R} \mathcal{A}(\phi|_{\mathrm{dom}(\varphi)}) \ : \ x\phi =_{A_\eta} [\![\mathsf{h}_s]\!]_{A_\eta}(\widehat{T}, y\phi)]
\end{aligned}$$

In the last probability expression, observe that $x\phi$ is drawn from the distribution $(\xleftarrow{R} [\![\mathit{Hash}]\!]_{A_\eta})$ independently from $\widehat{T}$ and $y\phi$. Hence, as the distribution $(\xleftarrow{R} [\![\mathit{Hash}]\!]_{A_\eta})$ is collision-free, the advantage of $\mathcal{B}$ is non-negligible; contradiction.

3. Let T_1 and T_2 be two terms of sort s such that $T_1 =_E T_2$. Consider the same frame as before: $\varphi = \{x_1 = \mathsf{h}_s(T_1, k),\ x_2 = \mathsf{h}_s(T_2, k)\}$. As $T_1 =_E T_2$ and h_s is free, we have $\varphi \approx_E \{x_1 = n,\ x_2 = n\}$ where n is a fresh name of sort *Hash* (Proposition 19 of Appendix A). By assumption this entails that $[\![\varphi]\!] \approx [\![\{x_1 = n,\ x_2 = n\}]\!]$ thus

$$\mathbb{P}\left[e'_1, e'_2 \xleftarrow{R} [\![\mathsf{h}_s(T_1, k),\ \mathsf{h}_s(T_2, k)]\!]_{A_\eta} \ : \ e'_1 =_{A_\eta} e'_2 \right] \geq 1 - \epsilon_\eta$$

where ϵ_η is a negligible function. As the implementation of h_s is collision-resistant, we deduce that

$$\mathbb{P}\left[e_1, e_2 \xleftarrow{R} [\![T_1,\ T_2]\!]_{A_\eta} \ : \ e_1 \neq_{A_\eta} e_2 \right]$$

is negligible. Other properties follow from Proposition 1. □

3.2. $\approx_E$-soundness implies classical assumptions on groups

In this section we present some interesting consequences of $\approx_E$-soundness. Inspired by the work of Hohenberger and Rivest on pseudo-freeness [33, 34], we prove that several standard cryptographic assumptions on groups are implied by the soundness of static equivalence. We concentrate on abelian groups as these are more relevant for cryptographic applications. We believe that similar techniques would apply for non-commutative groups as well.

We model an abelian group G with exponents taken over a commutative ring A by an abstract algebra over the following signature:

$$\begin{array}{lcllcl}
* & : & G \times G \to G & - & : & A \to A \\
1_G & : & G & \cdot & : & A \times A \to A \\
+ & : & A \times A \to A & 1_A & : & A \\
0 & : & A & \mathsf{exp} & : & G \times A \to G
\end{array}$$

We use the infix notation for the operators $*$, $\cdot$, $+$, and write g^a to denote $\mathsf{exp}(g, a)$. Note that the inverse operation on G is represented here by $g \mapsto \mathsf{exp}(g, -(1_A)) = g^{-(1_A)}$. We consider the equational theory E_{G} generated by the following equations (where x, y, z are variables of sort G, and u, v, w variables of sort A):

$$\begin{array}{rcl@{\qquad}rcl} u+v & = & v+u & x*y & = & y*x \\ u+(v+w) & = & (u+v)+w & x*(y*z) & = & (x*y)*z \\ u+0_A & = & u & x*1_G & = & x \\ u+(-u) & = & 0_A & (x^u)^v & = & x^{(u\cdot v)} \\ u\cdot v & = & v\cdot u & x^u * x^v & = & x^{u+v} \\ u\cdot(v\cdot w) & = & (u\cdot v)\cdot w & x^{1_A} & = & x \\ u\cdot 1_A & = & u & x^{0_A} & = & 1_G \\ (u+v)\cdot w & = & u\cdot w + v\cdot w & (x*y)^u & = & x^u * y^u \end{array}$$

We now recall several classical problems on groups. For cryptographic applications, it is desirable that these problems be *hard*, that is, not feasible by any probabilistic polynomial-time adversary:

- *discrete logarithm* (DL) problem: given g and g', find a, such that $g^a = g'$;
- *computational Diffie-Hellman* (CDH) problem: given g, g^a and g^b, find g^{ab};
- *decisional Diffie-Hellman* (DDH) problem: given g, g^a and g^b, distinguish g^{ab} from a random element g^c;
- *RSA* problem: given elements a and g^a, find g.

A more detailed presentation of these hard problems can be found in [35].

Assume a family of computational algebras (A_η) over the signature above such that (A_η) is $\approx_{E_{\mathsf{G}}}$-sound, at least for some subset of well-formed frames WF. Consider the two frames

$$\varphi_1 = \nu g, a, b.\{x_1 = g, x_2 = g^a, x_3 = g^b, x_4 = g^{a\cdot b}\} \text{ and}$$
$$\varphi_2 = \nu g, a, b, c.\{x_1 = g, x_2 = g^a, x_3 = g^b, x_4 = g^c\}.$$

and assume that φ_1, $\varphi_2 \in WF$. Then no probabilistic polynomial-time adversary $\mathcal{A}$ can solve the DDH problem in (A_η) with non-negligible probability.

Indeed, as suggested in [7], the question of (computationally) distinguishing these two frames exactly encodes the DDH problem. Given the equational theory E_{G}, we prove the formal equivalence $\varphi_1 \approx_{E_{\mathsf{G}}} \varphi_2$ (Lemma 21 of Appendix B). Thus, by $\approx_{E_{\mathsf{G}}}$-soundness, the DDH problem is hard in (A_η).

Clearly, if one can solve the DL problem, one can also solve the CDH problem, which itself allows us to solve the DDH problem. Therefore, the hardness of DDH implies the hardness of the two other problems.

In a similar way, we see that $\approx_{E_{\mathsf{G}}}$-soundness on an augmented signature implies the hardness of RSA. Instead of directly encoding the RSA problem, we introduce a slightly weaker decision problem, whose hardness implies the

hardness of RSA. The encoding of this problem requires the extension of the signature by a unary function symbol $\mathsf{h} : G \to \mathit{Hash}$, adding no equation to the theory. Consider the two frames

$$\varphi_3 = \nu g, a.\{x_1 = g^a, x_2 = a, x_3 = \mathsf{h}(g)\} \text{ and}$$
$$\varphi_4 = \nu g, a, h.\{x_1 = g^a, x_2 = a, x_3 = h\}.$$

We prove that $\varphi_3 \approx_{E_\mathsf{G}} \varphi_4$ in Lemma 22 of Appendix B. As above, if an implementation (A_η) is $\approx_{E_\mathsf{G}}$-sound of for some subset of well-formed frames *WF* including φ_3 and φ_4, then the RSA problem cannot be efficiently solved in (A_η). Indeed, any adversary $\mathcal{A}$ to the RSA-problem can be turned to an (equally efficient) adversary against $(\llbracket\varphi_3\rrbracket_{A_\eta}) \approx (\llbracket\varphi_4\rrbracket_{A_\eta})$ simply as follows: given a sample $\{x_1 = e_1, x_2 = e_2, x_3 = e_3\}$ from either side, let e be the result of $\mathcal{A}$ applied on η, e_1 and e_2; return 1 ("left-hand side") if $\llbracket h\rrbracket_{A_\eta}(e)$ equals to e_3, 0 otherwise.

An interesting open question is whether $\approx_{E_\mathsf{G}}$-soundness implies or is implied by Rivest's notion of pseudo-free groups [34], or equivalently [36], the strong RSA property. We conjecture that the two notions are in fact incomparable. Indeed, on the one hand, our notion implies the hardness of DDH, which remains an open question for strong RSA. On the other hand, pseudo-freeness and strong RSA deal with a form of adaptive attackers while our model is purely non-adaptive.

4. A sufficient (and often necessary) criterion for $\approx_E$-soundness

We now present useful results for proving $\approx_E$-soundness properties in general. Notably, we provide a sufficient criterion for $\approx_E$-soundness in Section 4.1 and prove it necessary under additional assumptions in Section 4.2.

4.1. Ideal semantics and $\approx_E$-soundness criterion

Given an implementation of the primitives, we have defined in Section 2.3 the concrete semantics $\llbracket\varphi\rrbracket_{A_\eta}$ associated to every frame φ . We now define the *ideal semantics* of a frame φ, intuitively as the conditional distribution over all the concrete values (in the appropriate space) that pass every formal test satisfied by φ.

Specifically, for every frame φ, we define the *tests* of φ to be

$$\mathrm{test}(\varphi) = \{(M, N) \mid \mathrm{var}(M, N) \subseteq \mathrm{dom}(\varphi),\ \mathrm{names}(M, N) \cap \mathrm{names}(\varphi) = \emptyset\}.$$

We let $\mathrm{eq}_E(\varphi)$ be the set of tests that are true in φ:

$$\mathrm{eq}_E(\varphi) = \{(M, N) \in \mathrm{test}(\varphi) \mid M\varphi =_E N\varphi\}$$

Note that, by definition, $\varphi \approx_E \varphi'$ iff $\mathrm{eq}_E(\varphi) \cap \mathrm{test}(\varphi') = \mathrm{eq}_E(\varphi') \cap \mathrm{test}(\varphi)$.

Let (A_η) be a family of computational algebras, $\varphi = \{x_1 = T_1, \ldots, x_n = T_n\}$ be a frame, and s_i be the sort of x_i. We define the set of eligible, well-formed values for φ by

$$\mathrm{Val}_{A_\eta}(\varphi) = \left\{\{x_1 = e_1, \ldots, x_n = e_n\} \mid (e_1, \ldots, e_n) \in \llbracket s_1\rrbracket_{A_\eta} \times \cdots \times \llbracket s_n\rrbracket_{A_\eta}\right\}$$

and write $\phi \xleftarrow{R} [\![\varphi]\!]^{\mathrm{val}}_{A_\eta}$ for the process of drawing a random value $\phi = \{x_1 = e_1, \ldots, x_n = e_n\}$ from $\mathrm{Val}_{A_\eta}(\varphi)$ using the drawings $e_i \xleftarrow{R} [\![s_i]\!]_{A_\eta}$ in the natural way.

Consider the following subset of concrete frames, intuitively, that pass all the valid tests of φ:

$$\mathrm{Val}'_{A_\eta}(\varphi) = \Big\{\phi \in \mathrm{Val}_{A_\eta}(\varphi) \mid \forall (M,N) \in \mathrm{eq}_E(\varphi), \\ \mathbb{P}\left[u, v \xleftarrow{R} [\![M,N]\!]_{A_\eta, \{x_1=e_1,\ldots,x_n=e_n\}} \ : \ u = v\right] = 1\Big\}$$

Note that, provided that (A_η) is unconditionally $=_E$-sound, $\mathrm{Val}'_{A_\eta}(\varphi)$ is non-empty as it contains at least the values given by the usual semantics of φ.

Definition 3 (Ideal semantics). Let (A_η) be an unconditionally $=_E$-sound family of computational algebras and φ be a frame. The *ideal semantics* of φ is the family of the distributions $[\![\varphi]\!]^{\mathrm{ideal}}_{A_\eta}$ obtained by conditionning each distribution $[\![\varphi]\!]^{\mathrm{val}}_{A_\eta}$ to the set of values $\mathrm{Val}'_{A_\eta}(\varphi)$. In other words, the probability to draw $\phi \in \mathrm{Val}_{A_\eta}(\varphi)$ is

$$\mathbb{P}[\phi \leftarrow [\![\varphi]\!]^{\mathrm{ideal}}_{A_\eta}] = \begin{cases} 0 & \text{if } \phi \notin \mathrm{Val}'_{A_\eta}(\varphi) \\ \frac{1}{V}\,\mathbb{P}[\phi \xleftarrow{R} [\![\varphi]\!]^{\mathrm{val}}_{A_\eta}] & \text{otherwise} \end{cases}$$

where $V = \mathbb{P}[\phi_0 \xleftarrow{R} [\![\varphi]\!]^{\mathrm{val}}_{A_\eta} \ : \ \phi_0 \in \mathrm{Val}'(\varphi)]$.

We say that (A_η) *has uniform distributions* if and only if for every η and every sort s, $[\![s]\!]_{A_\eta}$ is a finite set, $=_{A_\eta,s}$ is the usual equality, and the distribution associated to s by A_η is the uniform one over $[\![s]\!]_{A_\eta}$.

By classical property of conditional probabilities, we note that in the case of uniform distributions, the ideal semantics of a frame φ coincides with the family of uniform distributions over the (finite, non-empty) sets $\mathrm{Val}'_{A_\eta}(\varphi)$.

For instance, let $\varphi = \nu n_1, n_2.\{x_1 = n_1, x_2 = n_2\}$ with n_1 and n_2 of sort s. Then, given that E is stable by substitution of names, we have that $\mathrm{eq}_E(\varphi) = \{(M,N) \in \mathrm{test}(\varphi) \mid M =_E N\}$. By unconditional $=_E$-soundness, we deduce that $[\![\varphi]\!]^{\mathrm{ideal}}_{A_\eta}$ is simply the uniform distribution over $[\![s]\!]_{A_\eta} \times [\![s]\!]_{A_\eta}$.

We now state our $\approx_E$-soundness criterion: intuitively, the two semantics, concrete and ideal, should be indistinguishable.

Proposition 3 ($\approx_E$-soundness criterion). *Let (A_η) be an unconditionally $=_E$-sound family of computational algebras. Assume that for every frame φ it holds that $([\![\varphi]\!]_{A_\eta}) \approx ([\![\varphi]\!]^{\mathrm{ideal}}_{A_\eta})$. Then (A_η) is $\approx_E$-sound.*

Proof. Let $\varphi_1 \approx_E \varphi_2$. The equality $\mathrm{eq}_E(\varphi_1) \cap \mathrm{test}(\varphi_2) = \mathrm{eq}_E(\varphi_2) \cap \mathrm{test}(\varphi_1)$ entails $\mathrm{Val}'_{A_\eta}(\varphi_1) = \mathrm{Val}'_{A_\eta}(\varphi_2)$, thus the distributions $[\![\varphi_1]\!]^{\mathrm{ideal}}_{A_\eta}$ and $[\![\varphi_2]\!]^{\mathrm{ideal}}_{A_\eta}$ are equal. We use the transitivity of the indistinguishability relation $\approx$ to conclude: $([\![\varphi_1]\!]_{A_\eta}) \approx ([\![\varphi_1]\!]^{\mathrm{ideal}}_{A_\eta}) = ([\![\varphi_2]\!]^{\mathrm{ideal}}_{A_\eta}) \approx ([\![\varphi_2]\!]_{A_\eta})$. □

4.2. Transparent frames

In this section we show that our soundness criterion is necessary for a general class of equational theories, called *transparent* theories. In those theories, each frame can be associated to an equivalent *transparent frame* (defined below), which is easier to analyze.

Definition 4 (Transparent frames). A frame φ is *transparent* for an equational theory E if each of its subterms is deducible from φ in E.

Example 1. In the theory E_{enc}, the frame $\varphi_1 = \{x_1 = \mathsf{enc}(\mathsf{enc}(k_4, k_3), k_1),\ x_2 = \mathsf{enc}(k_1, k_2),\ x_3 = k_2\}$ is not transparent, as neither k_3 nor k_4 are deducible, but the frame $\overline{\varphi_1} = \{x_1 = \mathsf{enc}(n_1, k_1), x_2 = \mathsf{enc}(k_1, k_2), x_3 = k_2\}$ is.

The following proposition finitely characterizes the equations verified by a transparent frame.

Proposition 4. *Let φ be a transparent frame for E. Then, φ is of the form*

$$\varphi = \{x_1 = C_1[a_1, \ldots, a_m], \ldots, x_n = C_n[a_1, \ldots, a_m]\}$$

where $C_1, \ldots, C_n$ are (not necessarily linear) contexts such that $\mathrm{names}(C_1, \ldots, C_n) = \emptyset$*, $C_1[a_1, \ldots, a_m], \ldots, C_n[a_1, \ldots, a_m]$ are closed and, $a_1, \ldots, a_m$ are distinct deducible names: $\varphi \vdash_E a_i$.*

For each a_i, let M_{a_i} be a term such that $\mathrm{var}(M_{a_i}) \subseteq \{x_1, \ldots, x_n\}$, $\mathrm{names}(M_{a_i}) \cap \mathrm{names}(\varphi) = \emptyset$ *and $M_{a_i}\varphi =_E a_i$. Then every equation which holds in φ is a logical consequence of E and the equations $x_j = C_j[M_{a_1}, \ldots, M_{a_m}]$, written*

$$E \cup \{x_j = C_j[M_{a_1}, \ldots, M_{a_m}] \mid 1 \leq j \leq n\} \models \mathrm{eq}_E(\varphi).$$

By logical consequence, we refer to the usual first-order theory of equality, where the variables $x_1, \ldots, x_n$ are considered here as constants.

Proof. Let $(M, N) \in \mathrm{eq}_E(\varphi)$. By definition, we have $M\varphi =_E N\varphi$, that is, $M\{x_j \mapsto C_j[a_1, \ldots, a_m]\}_{1 \leq j \leq n} =_E N\{x_j \mapsto C_j[a_1, \ldots, a_m]\}_{1 \leq j \leq n}$. Since E is stable by substitution of names, we obtain

$$M\{x_j \mapsto C_j[M_{a_1}, \ldots, M_{a_m}]\}_{1 \leq j \leq n} =_E N\{x_j \mapsto C_j[M_{a_1}, \ldots, M_{a_m}]\}_{1 \leq j \leq n}.$$

Using the equalities $x_j = C_j[M_{a_1}, \ldots, M_{a_m}]$ and by transitivity, we obtain $\{x_j = C_j[M_{a_1}, \ldots, M_{a_m}] \mid 1 \leq j \leq n\} \cup E \models M = N$. □

Another nice and useful property of transparent frames is that their concrete and ideal semantics coincide.

Proposition 5. *Let (A_η) be an unconditionally $=_E$-sound family of computational algebras, having uniform distributions. Let φ be a transparent frame. The concrete and the ideal semantics of φ yield the same family of distributions: for all η, $[\![\varphi]\!]_{A_\eta} = [\![\varphi]\!]^{\mathrm{ideal}}_{A_\eta}$.*

PROOF. Let $\varphi = \{x_1 = C_1[a_1, \ldots, a_m], \ldots, x_n = C_n[a_1, \ldots, a_m]\}$, with $M_i\varphi =_E a_i$ $(1 \leq i \leq m)$ as above. Let s_i be the sort of a_i, s'_j be the sort of x_j and η a given complexity parameter.

The usual concrete semantics of φ consists in mapping every drawing of names from the set $\mathcal{E} = [\![s_1]\!]_{A_\eta} \times \cdots \times [\![s_m]\!]_{A_\eta}$ to a value in $\mathcal{F} = \mathrm{Val}_{A_\eta}(\varphi)$. Let us note $\alpha : \mathcal{E} \to \mathcal{F}$ this function, defined by:

$$\alpha(e_1, \ldots, e_m) = \Big\{ x_1 = [\![C_1[y_1, \ldots, y_m]]\!]_{\{y_1 = e_1, \ldots, y_m = e_m\}}, \ldots, \\ x_n = [\![C_n[y_1, \ldots, y_m]]\!]_{\{y_1 = e_1, \ldots, y_m = e_m\}} \Big\}$$

where the y_i are fresh variables respectively of sort s_i, and we omit the subscript A_η for sake of clarity.

Using the M_i, we can also define a function $\beta : \mathcal{F} \to \mathcal{E}$:

$$\beta(\phi) = \Big([\![M_1]\!]_\phi, \ldots, [\![M_m]\!]_\phi \Big)$$

We note that the distribution of $[\![M_i\varphi]\!]$ equals to that of $[\![M_i]\!]_\phi$ where $\phi \xleftarrow{R} [\![\varphi]\!]$, or equivalently, of $[\![M_i]\!]_{\alpha(e_1, \ldots, e_n)}$ where $(e_1, \ldots, e_n) \xleftarrow{R} \mathcal{E}$. As $M_i\varphi =_E a_i$, (A_η) is unconditionally $=_E$-sound, and no element of $\mathcal{E}$ has probability 0, we obtain that $\beta \circ \alpha = Id_E$. Thus α is injective and yields a bijection from $\mathcal{E}$ to its image $\mathcal{G} = \alpha(\mathcal{E})$. By assumption, $\mathcal{E}$ is equipped with the uniform distribution, therefore the concrete semantics of φ is the uniform distribution on $\mathcal{G}$.

Moreover $\mathcal{G}$ satisfies:

$$\begin{aligned} \mathcal{G} &= \{\phi \in \mathcal{F} \mid \alpha(\beta(\phi)) = \phi\} \\ &= \Big\{ \phi \in \mathcal{F} \mid \forall j, [\![C_j[y_1, \ldots, y_m]]\!]_{\{y_1 = [\![M_1]\!]_\phi, \ldots, y_m = [\![M_i]\!]_\phi\}} = [\![x_j]\!]_\phi \Big\} \\ &= \Big\{ \phi \in \mathcal{F} \mid \forall j, [\![C_j[M_1, \ldots, M_m]]\!]_\phi = [\![x_j]\!]_\phi \Big\} \end{aligned}$$

As φ is transparent, by Proposition 4, $\mathrm{eq}_E(\varphi)$ is implied by the equations $C_j[M_1, \ldots, M_m] = x_j$ and E. By unconditional $=_E$-soundness, we deduce that the values in $\mathcal{G}$ pass all the tests in $\mathrm{eq}_E(\varphi)$; in other words, $\mathcal{G} \subseteq \mathrm{Val}'_{A_\eta}(\varphi)$. Conversely, every element of $\mathrm{Val}'_{A_\eta}(\varphi)$ is trivially in $\mathcal{G}$; therefore $\mathcal{G} = \mathrm{Val}'_{A_\eta}(\varphi)$. Since $\mathcal{F}$ is equipped with uniform distribution, we obtain that the ideal semantics of φ coincides with the uniform distribution on $\mathcal{G}$, and therefore with its concrete semantics. □

A noticeable consequence of Proposition 5 is that, in the case of uniform distributions, two statically-equivalent transparent frames are always indistinguishable. (The argument is similar to that of Proposition 3.) This motivates the following definition, for the purpose of studying $\approx_E$-soundness or a converse to Proposition 3.

Definition 5. An equational theory E is *transparent* if and only if for every frame φ, there exists a (not necessarily unique) transparent frame $\overline{\varphi}$ such that $\varphi \approx_E \overline{\varphi}$.

Transparent frames and theories are related to the notion of *patterns* introduced by Abadi and Rogaway [4] and used in subsequent work [13, 12] so as to define computationally sound formal equivalences. There, messages are first mapped to patterns by replacing non-deducible subterms with boxes $\Box$. By definition, two messages are then equivalent if and only if they yield the same pattern (up to renaming of names). For example, if $\{M\}_K$ denotes the probabilistic encryption of M by a key K, the message $(\{\{K_4\}_{K_3}\}_{K_1}, \{K_1\}_{K_2}, K_2)$ is mapped to the pattern $(\{\Box\}_{K_1}, \{K_1\}_{K_2}, K_2)$. (Compare with example 1 where we have $\varphi_1 \approx_{E_{\mathsf{enc}}} \overline{\varphi_1}$.)

However, the notion of transparent frames is defined for any equational theory. Also, it might be the case that a frame corresponds to several transparent frames. For example, consider the theory of the exclusive OR (given in Section 5.1) and the frame:

$$\varphi = \{x_1 = n_1 \oplus n_2,\ x_2 = n_2 \oplus n_3,\ x_3 = n_1 \oplus n_3\}.$$

There are several transparent frames equivalent to φ, for instance $\{x_1 = n_1 \oplus n_2,\ x_2 = n_1,\ x_3 = n_2\}$, $\{x_1 = n_1,\ x_2 = n_1 \oplus n_2,\ x_3 = n_2\}$ and $\{x_1 = n_1,\ x_2 = n_2,\ x_3 = n_1 \oplus n_2\}$.

We believe that the notion of transparent frames is relevant in many theories useful in cryptography. As a matter of fact, the two theories of exclusive OR and ciphers considered in Section 5 are transparent. However, the notion of transparent frames does not subsume that of patterns, defined by Abadi and Rogaway. In particular, for the theory of probabilistic symmetric encryption, that is,

$$E_{\mathsf{senc}} = \{\mathsf{sdec}(\mathsf{senc}(x, y, z), y) = x, \quad \mathsf{sdec_success}(\mathsf{senc}(x, y, z), y) = \mathsf{ok}\},$$

it is unclear how to associate an equivalent transparent frame to the frame $\nu n, k, r.\{x = \mathsf{senc}(n, k, r), y = k\}$, although it is arguably a pattern in the sense of Abadi and Rogaway (once cast into our syntax). The reason is that the random coin r is not deducible, but the term $\mathsf{senc}(n, k, r)$ cannot be replaced with a fresh name because of the visible equation $\mathsf{sdec_success}(x, y) = \mathsf{ok}$. We might exclude r from being a subterm by modifying the notion of subterms (for example, in Abadi and Rogaway's work, the random factor does not appear explicitely in terms). However, this would undermine the properties of transparent frames mentioned above. Thus, we regard the notions of patterns and transparent frames as complementary.

Note that we have proved *en passant* that $\approx_E$ is decidable for transparent theories E for which $=_E$ is decidable, provided that the reduction to equivalent transparent frames is effective. Indeed, given two frames φ_1 and φ_2, we associate to each of them one of its statically equivalent transparent frame $\overline{\varphi_1}$ and $\overline{\varphi_2}$, respectively. It is then straightforward to check whether $\overline{\varphi_1}$ and $\overline{\varphi_2}$ are equivalent using the finite characterization of $\mathrm{eq}_E(\overline{\varphi_i})$ by Proposition 4.

Finally, we establish a completeness result for our soundness criterion in the cases of transparent theories.

Theorem 6. *Assume a transparent theory E. Let (A_η) be a family of computational algebras such that (A_η) has uniform distributions, is $\approx_E$-sound and unconditionally $=_E$-sound. Then the soundness criterion of Proposition 3 is satisfied: for every frame φ, $(\llbracket\varphi\rrbracket_{A_\eta}) \approx (\llbracket\varphi\rrbracket^{\mathrm{ideal}}_{A_\eta})$.*

Proof. Since E is transparent, there exists a transparent frame $\overline{\varphi}$ such that $\varphi \approx_E \overline{\varphi}$. By $\approx_E$-soundness, we deduce $(\llbracket\varphi\rrbracket_{A_\eta}) \approx (\llbracket\overline{\varphi}\rrbracket_{A_\eta})$. By Proposition 5, we have that $(\llbracket\overline{\varphi}\rrbracket_{A_\eta}) = (\llbracket\overline{\varphi}\rrbracket^{\mathrm{ideal}}_{A_\eta})$. Altogether, we conclude that $(\llbracket\varphi\rrbracket_{A_\eta}) \approx (\llbracket\varphi\rrbracket^{\mathrm{ideal}}_{A_\eta})$ since $\varphi \approx_E \overline{\varphi}$ implies $(\llbracket\overline{\varphi}\rrbracket^{\mathrm{ideal}}_{A_\eta}) = (\llbracket\varphi\rrbracket^{\mathrm{ideal}}_{A_\eta})$ as before. □

5. Examples

We now apply the framework of Sections 3 and 4 to establish two $\approx_E$-soundness results, concerning the theory of exclusive OR and that of ciphers and lists.

5.1. Exclusive OR

We study the soundness and faithfulness problems for the natural theory and implementation of the exclusive OR (XOR), together with constants and (pure) random numbers.

The formal model consists of a single sort *Data*, an infinite number of names, the infix symbol $\oplus : \mathit{Data} \times \mathit{Data} \to \mathit{Data}$ and two constants $0, 1 : \mathit{Data}$. Terms are equipped with the equational theory $E_\oplus$ generated by:

$$\begin{array}{rclcrcl} (x \oplus y) \oplus z & = & x \oplus (y \oplus z) & \qquad & x \oplus x & = & 0 \\ x \oplus y & = & y \oplus x & & x \oplus 0 & = & x \end{array}$$

As an implementation, we define the computational algebras A_η, $\eta \geq 0$:

- the concrete domain $\llbracket\mathit{Data}\rrbracket_{A_\eta}$ is the set of bit-strings of length η, $\{0,1\}^\eta$, equipped with the uniform distribution;
- $\oplus$ is interpreted by the usual XOR function over $\{0,1\}^\eta$;
- $\llbracket 0\rrbracket_{A_\eta} = 0^\eta$ and $\llbracket 1\rrbracket_{A_\eta} = 1^\eta$.

In this setting, statically equivalent frames enjoy an algebraic characterization. Let AC be the equational theory corresponding to the two left-hand equations for associativity and commutativity. We use the other two equations as a rewriting system $\mathcal{R}_\oplus$

$$\begin{array}{rcl} x \oplus x & \to & 0 \\ x \oplus 0 & \to & x \end{array}$$

where we allow arbitrary AC-manipulations before and after each rewriting step. It is easy to show that $\mathcal{R}_\oplus$ is AC-convergent. Specifically, a term T is in $\mathcal{R}_\oplus/AC$-normal form (or simply *normal form* in the following) if and only if each name, variable and constant 1 occur at most once in T, and 0 does not occur in T unless $T = 0$.

Let $a_1, \ldots, a_n$ be distinct names. Using the rewriting system $\mathcal{R}_\oplus/AC$, every closed term T with $\mathrm{names}(T) \subseteq \{a_1, \ldots, a_n\}$ can be written $T =_{E_\oplus} \beta_0 \oplus \bigoplus_{j=1}^n \beta_j\, a_j$ where $\beta_j \in \{0,1\}$, the a_j are mutually distinct, and we use the convention $0a_j = 0$ and $1a_j = a_j$. In the following, we see $\{0,1\}$ as the two-element field $\mathbb{F}_2$; thus terms modulo $=_{E_\oplus}$ form a $\mathbb{F}_2$-vector space.

Similarly a frame φ with $\mathrm{names}(\varphi) \subseteq \{a_1, \ldots, a_n\}$ is written

$$\varphi =_{E_\oplus} \left\{ x_1 = \alpha_{1,0} \oplus \bigoplus_{j=1}^n \alpha_{1,j}\, a_j \ ,\ \ldots\ ,\ x_m = \alpha_{m,0} \oplus \bigoplus_{j=1}^n \alpha_{m,j}\, a_j \right\}$$

where $\alpha_{i,j} \in \mathbb{F}_2$. Let us group the coefficients into a $(m+1) \times (n+1)$-matrix $\alpha = (\alpha_{i,j})$ over $\mathbb{F}_2$. Then, φ is described by the formal relation

$$\begin{pmatrix} 1 \\ x_1 \\ \vdots \\ x_m \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 0 & \ldots & 0 \\ \alpha_{1,0} & \alpha_{1,1} & \ldots & \alpha_{1,n} \\ \vdots & & & \vdots \\ \alpha_{m,0} & \alpha_{m,1} & \ldots & \alpha_{m,n} \end{pmatrix}}_{\alpha} \cdot \begin{pmatrix} 1 \\ a_1 \\ \vdots \\ a_n \end{pmatrix}$$

We now characterize the set $\mathrm{eq}_{E_\oplus}(\varphi)$ of equations valid in φ. Let M and N be two terms such that $\mathrm{var}(M,N) \subseteq \mathrm{dom}(\phi), \mathrm{names}(M,N) \cap \mathrm{names}(\varphi) = \emptyset$. First note that $M\varphi =_{E_\oplus} N\varphi$ if and only if $(M \oplus N)\varphi =_{E_\oplus} 0$. Therefore we only study the case where $N = 0$.

Assume M in normal form. $M\varphi =_{E_\oplus} 0$ and $\mathrm{names}(M) \cap \mathrm{names}(\varphi) = \emptyset$ implies $\mathrm{names}(M) = \emptyset$. Let $M =_{AC} \beta_0 \oplus \bigoplus_{i=1}^m \beta_i\, x_i$. The condition $M\varphi =_{E_\oplus} 0$ is equivalent to the vectorial equation

$$(\beta_0, \ldots, \beta_m) \cdot \alpha = 0$$

that is, $(\beta_0, \ldots, \beta_m)$ belongs to the co-kernel of α, noted $\mathrm{coker}(\alpha)$.

Finally let φ and φ' be two frames with $\mathrm{names}(\varphi, \varphi') \subseteq \{a_1, \ldots, a_n\}$ and $\mathrm{dom}(\varphi) = \mathrm{dom}(\varphi') = \{x_1, \ldots, x_m\}$. Let α and α' be the two corresponding $(m+1) \times (n+1)$-matrices defined as above. From the previous discussion, we deduce that

$$\varphi \approx_{E_\oplus} \varphi' \ \Leftrightarrow\ \mathrm{coker}(\alpha) = \mathrm{coker}(\alpha')$$

that is, if we write $\mathrm{im}(\alpha) = \{\alpha \cdot \gamma\}$ the image of α, we have by duality

$$\varphi \approx_{E_\oplus} \varphi' \ \Leftrightarrow\ \mathrm{im}(\alpha) = \mathrm{im}(\alpha'). \tag{1}$$

This characterization is the key point of our main result for the theory of XOR.

Theorem 7. *The implementation of XOR for the considered signature, (A_η), is unconditionally $=_{E_\oplus}$-, $\approx_{E_\oplus}$- and $\not\approx_{E_\oplus}$-sound. It is also $=_{E_\oplus}$-, $\approx_{E_\oplus}$- and $\not\approx_{E_\oplus}$-faithful.*

PROOF. The unconditional $=_{E_\oplus}$-soundness is clear, hence the $\not\models_{E_\oplus}$-faithfulness (Proposition 1).

Let us show that (A_η) is $=_{E_\oplus}$-faithful. Assume that T_1 and T_2 are two terms such that $T_1 \neq_{E_\oplus} T_2$. This is equivalent to $T_1 \oplus T_2 \neq_{E_\oplus} 0$. Thus it is sufficient to consider the case where $T \neq 0$ is a closed term in normal form. The semantics of T is either the constant 1^η (if $T = 1$) or the uniform distribution (if $T \neq 1$) on $\{0,1\}^\eta$. Thus $\mathbb{P}\left[[\![T]\!]_{A_\eta} = 0\right]$ is negligible. Hence the $=_{E_\oplus}$-faithfulness holds and by proposition 1, so does the $\approx_{E_\oplus}$-faithfulness.

We now address the unconditional $\approx_{E_\oplus}$-soundness. Let φ be a frame, and $\alpha = (\alpha_{i,j})$ its $(m+1) \times (n+1)$-matrix associated as before. Let us see α as a $\mathbb{F}_2$-linear function from $(\mathbb{F}_2)^{n+1}$ to $(\mathbb{F}_2)^{m+1}$.

For simplicity, let us fix the order of variables in $\mathrm{dom}(\varphi)$ and assimilate the possible concrete values of φ, $\mathrm{Val}_{A_\eta}(\varphi)$, to the set $\mathcal{F} = \{1^\eta\} \times (\mathbb{F}_2)^{m\eta}$ where the first η 1-bits are added for technical reasons.

The usual concrete semantics of φ consists in drawing a random vector uniformly from $\mathcal{E} = \{1^\eta\} \times (\mathbb{F}_2)^{n\eta}$ for the value of names, and then applying a $\mathbb{F}_2$-linear function $\widehat{\alpha} : (\mathbb{F}_2)^{(n+1)\eta} \to (\mathbb{F}_2)^{(m+1)\eta}$ to it. Specifically, if we see $(\mathbb{F}_2)^{(n+1)\eta}$ as $\underbrace{{\mathbb{F}_2}^\eta \times \ldots \times {\mathbb{F}_2}^\eta}_{n+1}$ and similarly for $(\mathbb{F}_2)^{(m+1)\eta}$, the function $\widehat{\alpha}$ is defined by

$$\widehat{\alpha}\,(f_0, \ldots, f_n) = \left(\bigoplus_{j=0}^{n} \alpha_{0,j}\, f_j, \ldots, \bigoplus_{j=0}^{n} \alpha_{m,j}\, f_j \right)$$

Since $\widehat{\alpha}$ is linear, all the inverse images $\widehat{\alpha}^{-1}(\{x\})$, $x \in \mathrm{im}(\widehat{\alpha})$, have the same cardinal. Hence, the concrete semantics of φ is also the uniform distribution over $\widehat{\alpha}(\mathcal{E}) = \mathrm{im}(\widehat{\alpha}) \cap \mathcal{F}$.

Assume a second frame φ' such that $\varphi \approx_{E_\oplus} \varphi'$. Define α' and $\widehat{\alpha'}$ similarly as above. By equation 1, we have $\mathrm{im}(\alpha) = \mathrm{im}(\alpha')$.

Now, if we see $(\mathbb{F}_2)^{(m+1)\eta}$ as $\underbrace{{\mathbb{F}_2}^{m+1} \times \ldots \times {\mathbb{F}_2}^{m+1}}_{\eta}$, we may write $\widehat{\alpha} = \underbrace{\alpha \times \ldots \times \alpha}_{\eta}$ and similarly for α'. Thus,

$$\mathrm{im}(\widehat{\alpha}) = \underbrace{\mathrm{im}(\alpha) \times \ldots \times \mathrm{im}(\alpha)}_{\eta} = \underbrace{\mathrm{im}(\alpha') \times \ldots \times \mathrm{im}(\alpha')}_{\eta} = \mathrm{im}(\widehat{\alpha'})$$

which implies that φ and φ' have the same concrete semantics. Thus $E_\oplus$ is unconditionally $\approx_{E_\oplus}$-sound.

Last, we prove the unconditional $\not\models_{E_\oplus}$-soundness. Let φ be a frame and T a term, both in normal form, such that $\varphi \not\models_{E_\oplus} T$ and $\mathrm{names}(T) \subseteq \mathrm{names}(\varphi) = \{a_1, \ldots, a_n\}$. Let α be associated to φ as before and $T =_{AC} \beta_0 \oplus \bigoplus_{j=1}^{n} \beta_j\, a_j$.

Let γ be the $(m+2) \times (n+1)$-matrix obtained by augmenting α with a last

row equal to $\beta = (\beta_0, \ldots, \beta_n)$:

$$\gamma = \begin{pmatrix} 1 & 0 & \ldots & 0 \\ \alpha_{1,0} & \alpha_{1,1} & \ldots & \alpha_{1,n} \\ \vdots & & & \vdots \\ \alpha_{m,0} & \alpha_{m,1} & \ldots & \alpha_{m,n} \\ \beta_0 & \beta_1 & \ldots & \beta_{n_0} \end{pmatrix}$$

Since $\varphi \not\vdash_{E_\oplus} T$, in particular there exists no M in normal form such that $\mathrm{names}(M) = \emptyset$ and $M\varphi =_{E_\oplus} T$. In other words, β is linearly independent from the other rows in the matrix γ above.

In particular, it is independent from the first row $(1, 0, \ldots, 0)$, that is, there exists $j \geq 1$ such that $\beta_j \neq 0$. We deduce that the distribution $(\overleftarrow{}^R [\![T]\!]_{A_\eta})$ is the uniform one over $\{0,1\}^\eta$, thus it is collision-free.

As for the first condition of unconditional $\not\vdash_E$-soundness, by a similar reasoning as before, we have that the concrete semantics of (φ, T) is the uniform distribution over the image of $\mathcal{E} = \{1^\eta\} \times (\mathbb{F}_2)^{n\eta}$ by $\widehat{\gamma}$ (defined similarly as $\widehat{\alpha}$ above). Let us see β a linear function from $(\mathbb{F}_2)^{n+1}$ to $\mathbb{F}_2$ and define $\widehat{\beta}$ as previously. Next we prove that the image $\widehat{\gamma}(\mathcal{E})$ is the cartesian product of the two sets $\widehat{\alpha}(\mathcal{E})$ and $\widehat{\beta}(\mathcal{E})$. It follows that the drawings for φ and T are independent.

The inclusion $\widehat{\gamma}(\mathcal{E}) \subseteq \widehat{\alpha}(\mathcal{E}) \times \widehat{\beta}(\mathcal{E})$ is trivial. As β is independent from the rows of α, there exists a vector $u \in (\mathbb{F}_2)^{n+1}$ such that $\beta(u) = 1$ and $\alpha(u) = 0$ (otherwise $\ker(\beta) \supseteq \ker(\alpha)$ implies $\beta \in \mathrm{coim}(\beta) \subseteq \mathrm{coim}(\alpha)$). Let $x, y \in \mathcal{E}$. We prove that there exists $z \in \mathcal{E}$ such that $\widehat{\alpha}(z) = \widehat{\alpha}(x) \in (\mathbb{F}_2)^{(m+1)\eta}$ and $\widehat{\beta}(z) = \widehat{\beta}(y) \in (\mathbb{F}_2)^\eta$.

Indeed, let us see $\mathcal{E}$ as $(\{1\} \times (\mathbb{F}_2)^n)^\eta$. Using the corresponding bases, let $x = (x_1, \ldots, x_\eta)$ and $y = (y_1, \ldots, y_\eta)$ with $x_i, y_i \in \{1\} \times (\mathbb{F}_2)^n$. We let $z_i = x_i + (\beta(y_i) - \beta(x_i)) \cdot u$ and $z = (z_1, \ldots, z_\eta)$. Thus, $\widehat{\alpha}(z) = (\alpha(z_1), \ldots, \alpha(z_\eta)) = (\alpha(x_1), \ldots, \alpha(x_\eta)) = \widehat{\alpha}(x)$ and $\widehat{\beta}(z) = (\beta(z_1), \ldots, \beta(z_\eta)) = (\beta(y_1), \ldots, \beta(y_\eta)) = \widehat{\beta}(y)$. Besides, $\alpha(u) = 0$ implies that the first coordinate of u is 0, thus the first coordinate of each z_i is 1, that is, $z \in \mathcal{E}$. □

We conclude this section by a proof that the $E_\oplus$ is transparent as announced in Section 4.

Proposition 8. *The equational theory $E_\oplus$ is transparent.*

Proof. Indeed, let φ be frame and α be its associated $(m+1) \times (n+1)$-matrix as before. Let d be the dimension of $\mathrm{im}(\alpha)$. There exists a $(m+1) \times d$ submatrix α' of α such that α' is injective and $\mathrm{im}(\alpha') = \mathrm{im}(\alpha)$ (consider a maximal independent set of columns of α). As the first column of α is independent from the others (it starts with a 1 whereas the others start with a 0), we may assume without loss of generality that the first column of α' is that of α. (In particular $d \geq 1$.)

Let $a'_1 \dots a'_{d-1}$ be distinct names. We let φ' be the frame associated to α', described by the relation

$$\begin{pmatrix} 1 \\ x_1 \\ \vdots \\ x_m \end{pmatrix} = \alpha' \cdot \begin{pmatrix} 1 \\ a'_1 \\ \vdots \\ a'_{d-1} \end{pmatrix}.$$

As $\mathrm{im}(\alpha') = \mathrm{im}(\alpha)$, we have $\varphi' \approx_{E_\oplus} \varphi$. Besides, since α' is injective, there exists α'' such that $\alpha'' \cdot \alpha'$ is the identity $d \times d$-matrix. This entails that every a'_i is deducible from φ', that is, φ' is transparent. □

5.2. Symmetric, deterministic, length-preserving encryption and lists

We now detail the example of symmetric, deterministic and length-preserving encryption schemes. Such schemes, also known as *pseudo-random permutations* or *ciphers* [37], are widely used in practice, the most famous examples (for fixed-length inputs) being DES and AES.

Our formal model consists of a set of sorts $\mathcal{S} = \{Data, List_0, List_1 \dots List_n \dots\}$, an infinite number of names for every sort $Data$ and $List_n$, and the following symbols (for every $n \geq 0$):

$\mathsf{enc}_n, \mathsf{dec}_n$	:	$List_n \times Data \to List_n$	encryption, decryption
cons_n	:	$Data \times List_n \to List_{n+1}$	list constructor
head_n	:	$List_{n+1} \to Data$	head of a list
tail_n	:	$List_{n+1} \to List_n$	tail of a list
nil	:	$List_0$	empty list
$0, 1$	:	$Data$	constants

We consider the equational theory E_{sym} generated by the following equations (for every $n \geq 0$ and for every name a_0 of sort $List_0$):

$$\begin{aligned} \mathsf{dec}_n(\mathsf{enc}_n(x,y),y) &= x \\ \mathsf{enc}_n(\mathsf{dec}_n(x,y),y) &= x \\ \mathsf{head}_n(\mathsf{cons}_n(x,y)) &= x \\ \mathsf{tail}_n(\mathsf{cons}_n(x,y)) &= y \\ \mathsf{cons}_n(\mathsf{head}_n(x),\mathsf{tail}_n(x)) &= x \end{aligned} \qquad \begin{aligned} \mathsf{enc}_0(\mathsf{nil},x) &= \mathsf{nil} \\ \mathsf{dec}_0(\mathsf{nil},x) &= \mathsf{nil} \\ \mathsf{tail}_0(x) &= \mathsf{nil} \\ a_0 &= \mathsf{nil} \end{aligned}$$

where x, y are variables of the appropriate sorts in each case. The effect of the last four equations is that the sort $List_0$ is degenerated in E_{sym}, that is, all terms of sort $List_0$ are equal. When oriented from left to right, the equations above form a convergent rewriting system written $\mathcal{R}$.

Notice that each term has a unique sort. As the subscripts n of function symbols are redundant with sorts, we tend to omit them in terms. For instance, if $k, k' : Data$, we may write $\mathsf{enc}(\mathsf{cons}(k, \mathsf{nil}), k')$ instead of $\mathsf{enc}_1(\mathsf{cons}_0(k, \mathsf{nil}), k')$.

The concrete meaning of sorts and symbols is given by the computational algebras A_η, $\eta > 0$, defined as follows:

- the carrier sets are $[\![Data]\!]_{A_\eta} = \{0,1\}^\eta$ and $[\![List_n]\!]_{A_\eta} = \{0,1\}^{n\eta}$ equipped with the uniform distribution and the usual equality relation;
- $\mathsf{enc}_n, \mathsf{dec}_n$ are implemented by a cipher for data of size $n\eta$ and keys of size η; (we discuss the required cryptographic assumptions later);
- $[\![\mathsf{nil}]\!]_{A_\eta}$ is the empty bit-string, $[\![\mathsf{cons}_n]\!]_{A_\eta}$ is the usual concatenation, $[\![0]\!]_{A_\eta} = 0^\eta$, $[\![1]\!]_{A_\eta} = 1^\eta$, $[\![\mathsf{head}_n]\!]_{A_\eta}$ returns the η first digits of bit-strings (of size $(n+1)\eta$) whereas $[\![\mathsf{tail}_n]\!]_{A_\eta}$ returns the last $n\eta$ digits.

We emphasize that no tags are added to messages. Tags—and in particular tags under encryption—would be harmful to the $\approx_{E_{\mathsf{sym}}}$-soundness. Indeed we expect that the formal equivalence $\nu a, b.\{x = \mathsf{enc}(a,b),\ y = b\} \approx_{E_{\mathsf{sym}}} \nu a, b, c.\{x = \mathsf{enc}(a,b),\ y = c\}$ also holds in the computational world; but this would not be the case if a is tagged before encryption. In case a was tagged before encryption, an adversary could use the tag to check the success of decrypting $\mathsf{enc}(a,b)$ with b.

For simplicity we assume without loss of generality that encryption keys have the same size η as blocks of data. We also assume that keys are generated according to the uniform distribution.

It is not difficult to prove that the above implementation is unconditionally $=_{E_{\mathsf{sym}}}$-sound (by induction on the structure of terms and equational proofs), that is, every true formal equality holds with probability 1 in the concrete world. We note that the equation $\mathsf{enc}_n(\mathsf{dec}_n(x,y),y) = x$ is satisfied because encryption by a given key is length-preserving and injective, hence also surjective.

Before studying the $\approx_{E_{\mathsf{sym}}}$-soundness, we need to characterize statically equivalent frames. Specifically, we show that this theory is transparent.

Proposition 9. *Let φ be a closed frame. There exists a transparent frame $\overline{\varphi}$ such that $\varphi \approx_{E_{\mathsf{sym}}} \overline{\varphi}$.*

The proof of Proposition 9 relies on the following Lemma 10, that is used stepwise to rewrite a frame into a transparent frame.

Lemma 10. *Let φ be a closed frame in $\mathcal{R}$-normal form. Let T be a subterm of φ of the form $T = \mathsf{enc}(U,V)$, $T = \mathsf{dec}(U,V)$, $T = \mathsf{head}(V)$ or, $T = \mathsf{tail}(V)$ and n a fresh name of the same sort than T. Assume that V is not deducible from φ, that is, $\varphi \not\vdash_{E_{\mathsf{sym}}} V$. Then we have that*

$$\varphi \approx_{E_{\mathsf{sym}}} \varphi'$$

where $\varphi' = \varphi\{T \mapsto n\}$ is obtained by replacing every *occurrence of T in φ with n.*

The proof of Lemma 10 is given in Appendix C. We prove Proposition 9 by applying this lemma repeatedly on an initial frame φ. The procedure terminates as each rewriting step decreases the total size of non-deducible subterms in the frame. Besides, the resulting frame $\overline{\varphi}$ is transparent. Indeed, by contradiction,

suppose that $\overline{\varphi}$ is not transparent; define T as the father of the largest non-deducible subterm of φ; it is easy to see that T is necessarily of the form $T = \mathsf{enc}(U, V)$, $T = \mathsf{dec}(U, V)$, $T = \mathsf{head}(V)$ or $T = \mathsf{tail}(V)$ with $\varphi \not\vdash_{E_{\mathsf{sym}}} V$; thus Lemma 10 applies.

Note that for any subterm W, $\varphi \not\vdash_{E_{\mathsf{sym}}} W$ implies $\varphi\{T \mapsto n\} \not\vdash_{E_{\mathsf{sym}}} W\{T \mapsto n\}$. As a consequence, the procedure above yields a unique transparent frame $\overline{\varphi}$ (modulo renaming), no matter in which order the subterms T are substituted.

Provided that $\vdash_{E_{\mathsf{sym}}}$ is decidable[1], the above procedure for associating transparent frames to frames is effective. Thus, as noticed in Section 4.2, we obtain another proof of the decidability of $\approx_{E_{\mathsf{sym}}}$ using Proposition 4. Notice that statically equivalent transparent frames may *not* be equal modulo renaming: consider for instance $\{x = \mathsf{enc}(a, b),\, y = b\} \approx_{E_{\mathsf{sym}}} \{x = c,\, y = b\}$.

We now study the $\approx_{E_{\mathsf{sym}}}$-soundness problem under classical cryptographic assumptions. Standard assumptions on ciphers include the notions of super pseudo-random permutation (SPRP) and several notions of indistinguishability (IND-Pi-Cj, $i, j = 0, 1, 2$). In particular, IND-P1-C1 denotes the indistinguishability against lunchtime chosen-plaintext and chosen-ciphertext attacks. These notions and the relations between them have been studied notably in [37].

Initially, the SPRP and IND-P1-C1 assumptions apply to (block) ciphers specialized to plaintexts of a given size. Interestingly, this is not sufficient to imply $\approx_{E_{\mathsf{sym}}}$-soundness for frames which contain plaintexts of heterogeneous sizes, encrypted under the same key. Thus we introduce a strengthened version of IND-P1-C1, applying to a *collection* of ciphers $(\mathcal{E}_{\eta,n}, \mathcal{D}_{\eta,n})$, where η is the complexity parameter and $n \geq 0$ is the number of blocks of size η contained in plaintexts and ciphertexts. One may note that there exist operation modes which turn a fixed size block cipher realizing SPRP into a cipher which handles variable length inputs while preserving SPRP. We refer the reader to [38] for an example of such a mode and further references.

We define the ω-IND-P1-C1 assumption by considering the following experiment $\mathcal{G}_\eta$ with a 2-stage adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$:

- first a key k is randomly chosen from $\{0,1\}^\eta$;
- (Stage 1) $\mathcal{A}_1$ is given access to the encryption oracles $\mathcal{E}_{\eta,n}(\cdot, k)$ and the decryption oracles $\mathcal{D}_{\eta,n}(\cdot, k)$; it outputs two plaintexts $m_0, m_1 \in \{0,1\}^{n_0\eta}$ for some n_0, and possibly some data d;
- (Stage 2) a random bit $b \in \{0,1\}$ is drawn; $\mathcal{A}_2$ receives the data d, the *challenge ciphertext* $c = \mathcal{E}_{\eta,n_0}(m_b, k)$ and outputs a bit b';
- $\mathcal{A}$ *is successful in* $\mathcal{G}_\eta$ iff $b = b'$ and it has never submitted m_0 or m_1 to an encryption oracle, nor c to a decryption oracle.

[1] A classical characterization of deducibility, entailing its decidability, is detailed in Lemma 23 of Appendix C.

Define the *advantage* of $\mathcal{A}$ as

$$\mathrm{Adv}_{\mathcal{A}}^{\omega\text{-IND-P1-C1}}(\eta) = 2 \times \mathbb{P}\left[\mathcal{A} \text{ is successful in } \mathcal{G}_\eta\right] - 1 \tag{2}$$

The ω-IND-P1-C1 assumption holds for $(\mathcal{E}_{\eta,n}, \mathcal{D}_{\eta,n})$ iff the advantage of any probabilistic polynomial-time adversary is negligible. It holds for the *inverse* of the encryption scheme iff it holds for the collection of ciphers $(\mathcal{D}_{\eta,n}, \mathcal{E}_{\eta,n})$.

As in previous work [4, 13, 18, 23], we restrict frames to those with only atomic keys and no encryption cycles. Specifically, a closed frame φ *has only atomic keys* if for all subterms $\mathsf{enc}_n(u, v)$ and $\mathsf{dec}_n(u, v)$ of φ, v is a name. Given two (atomic) keys k_1 and k_2, we say that k_1 *encrypts* k_2 *in* φ, written $k_1 >_\varphi k_2$, iff there exists a subterm U of φ of the form $U = \mathsf{enc}_n(T, k_1)$ or $U = \mathsf{dec}_n(T, k_1)$ such that k_2 appears in T *not used as a key*, that is, k_2 appears in T at a position which is not the right-hand argument of a $\mathsf{enc}_{n'}$ or a $\mathsf{dec}_{n'}$. An *encryption cycle* is a tuple $k_1 \dots k_m$ such that $k_1 >_\varphi \dots >_\varphi k_m >_\varphi k_1$.

The effect of the condition "not used as a key" is to allow considering more terms as free of encryption cycles, for instance $\mathsf{enc}_n(\mathsf{enc}_n(a, k), k)$. This improvement is already suggested in [4].

We now state our $\approx_{E_{\mathsf{sym}}}$-soundness theorem. A closed frame is *well-formed* iff its $\mathcal{R}$-normal form has only atomic keys, contains no encryption cycles and uses no head and tail symbols.

Theorem 11 ($\approx_{E_{\mathsf{sym}}}$-soundness). *Let φ_1 and φ_2 be two well-formed frames of the same domain. Assume that the concrete implementations for the encryption and its inverse satisfy both the ω-IND-P1-C1 assumption. If $\varphi_1 \approx_{E_{\mathsf{sym}}} \varphi_2$ then $(\llbracket\varphi_1\rrbracket_{A_\eta}) \approx (\llbracket\varphi_2\rrbracket_{A_\eta})$.*

Before proving Theorem 11, we establish a computational counterpart to Lemma 10.

Lemma 12. *Let φ be a closed frame in $\mathcal{R}$-normal form, with only atomic keys and no encryption cycles. Let T be a subterm of φ of the form $T = \mathsf{enc}(U, k)$ (respectively $T = \mathsf{dec}(U, k)$), with k name of sort Data, and n a fresh name of the same sort as T. Assume that*

- *the only occurrences of k in φ are in the positions of an encryption or decryption key: $\mathsf{enc}(., k)$ or $\mathsf{dec}(., k)$;*
- *T itself does not appear under an encryption or a decryption with k;*
- *the concrete implementations for the encryption and its inverse satisfy both the ω-IND-P1-C1 assumption.*

Then we have that

$$(\llbracket\varphi\rrbracket_{A_\eta}) \approx (\llbracket\varphi'\rrbracket_{A_\eta})$$

where $\varphi' = \varphi\{T \mapsto n\}$ is obtained by replacing every *occurrence of T in φ with n.*

Notice that the hypothesis of Lemma 12 are stronger than its formal version, Lemma 10. For instance the encryption key k is required to be atomic; the first condition on k implies that k is not deducible from φ. Also nothing is said about head and tail symbols.

Proof (of Lemma 12). Before proving the lemma, let us consider the example of a well-formed frame $\varphi_1 = \{x_1 = \mathsf{enc}(T_1, k),\ x_2 = \mathsf{enc}(T_2, k)\}$, where k does not appear in T_1, T_2, and $T_1 \neq_{E_{\mathsf{sym}}} T_2$. This frame is statically equivalent to $\varphi_2 = \{x_1 = n_1; x_2 = n_2\}$. Our problem here is to prove that $[\![\varphi_1]\!]$ and $[\![\varphi_2]\!]$ are actually indistinguishable. It is not hard to see that this will be the case if and only if the probability that T_1 and T_2 have the same concrete value is negligible. A consequence of this phenomenon is intuitively that we need to prove Lemma 12 and—at least—a limited form of $=_{E_{\mathsf{sym}}}$-faithfulness at the same time.

Formally, let us write $|\varphi|_e$ and $|T|_e$ for the number of distinct subterms with head symbols enc or dec, occurring respectively in a frame φ and a term T. Let P_n and Q_n be the two properties:

> (P_n) Lemma 12 holds provided that $|\varphi|_e \leq n$:
> For every $\mathcal{R}$-normal, closed frame φ containing only atomic keys, no encryption cycles, and such that $|\varphi|_e \leq n$, for every maximal subterm T of φ of the form $T = \mathsf{enc}(U, k)$ or $T = \mathsf{dec}(U, k)$, for every fresh name n of the appriopriate sort, if the only occurrences of k in φ are in key positions (*i.e.* $\mathsf{enc}(., k)$ or $\mathsf{dec}(., k)$), then $([\![\varphi]\!]_{A_\eta}) \approx ([\![\varphi\{T \mapsto n\}]\!]_{A_\eta})$.
>
> (Q_n) For all $\mathcal{R}$-normal terms T_1, T_2 of the same sort such that: T_1, T_2 have only atomic keys, the frame $\varphi = \{x = T_1, y = T_2\}$ has no encryption cycles, $T_1 \neq T_2$ and $|\varphi|_e \leq n$, the probability $\mathbb{P}\left[e_1, e_2 \leftarrow [\![T_1, T_2]\!]_{A_\eta}; e_1 = e_2\right]$ is negligible.

We prove P_n and Q_n by mutual induction on n, that is, more precisely we prove the four statements: (S1) P_0, (S2) $P_{n+1} \Leftarrow Q_n$, (S3) Q_0, (S4) $Q_{n+1} \Leftarrow (P_{n+1}$ and $Q_n)$.

(S1) P_0 is vacuously true.

(S2) $P_{n+1} \Leftarrow Q_n$. Let $T^0 = \mathsf{enc}_{n_0}(U, k)$ be a subterm of φ, k and n two names all satisfying the conditions of Lemma 12. (Naturally, the case of $T^0 = \mathsf{dec}_{n_0}(U, k)$ is similar.) Let $\varphi = \{x_1 = T_1^0, \ldots, x_n = T_n^0\}$.

Provided an adversary $\mathcal{A}$ able to distinguish $([\![\varphi]\!]_{A_\eta})$ and $([\![\varphi']\!]_{A_\eta})$, we build an adversary $\mathcal{B}$ against the ω-IND-P1-C1 assumption on encryption, described as follows:

1. for each name a of sort s appearing in φ, draw a value $\widehat{a} \xleftarrow{R} [\![s]\!]_{A_\eta}$;
2. draw a value $\widehat{a_0} \xleftarrow{R} [\![s]\!]_{A_\eta}$ for some fresh name a_0 of sort $List_{n_0}$;
3. for each x_i $(1 \leq i \leq n)$ of sort s_i, compute $\widehat{T_i^0} \in [\![s_i]\!]_A$ recursively as

follows:

$$\begin{aligned}
\widehat{\mathsf{enc}_n(T,k)} &= \mathcal{E}_n(\widehat{T}) \text{ if } T \neq U \\
\widehat{\mathsf{enc}_{n_0}(U,k)} &= \mathcal{E}^*(\widehat{U},\widehat{a_0}) \\
\widehat{\mathsf{dec}_n(T,k)} &= \mathcal{D}_n(\widehat{T}) \\
\widehat{f(T_1,\ldots,T_n)} &= [\![f]\!]_{A_\eta}(\widehat{T_1},\ldots,\widehat{T_n}) \quad \text{in the remaining cases}
\end{aligned}$$

where we have written $\mathcal{E}_n(.)$ and $\mathcal{D}_n(.)$ for the encryption and decryption oracles of the ω-IND-P1-C1 game, and $\mathcal{E}^*(\widehat{U},\widehat{a_0})$ for the challenge ciphertext, obtained after submitting the two plaintexts $\widehat{U}$ and $\widehat{a_0}$. Since $T^0 = \mathsf{enc}_{n_0}(U,k)$ is not a subterm of an encryption or a decryption with k, we may assume that $\mathcal{E}^*(\widehat{U},\widehat{a_0})$ is computed only once, after every call to $\mathcal{E}_n(.)$ and $\mathcal{D}_n(.)$;

4. submit the concrete frame $\{x_1 = \widehat{T_1}, \ldots, x_n = \widehat{T_n}\}$ to $\mathcal{A}$ and return the same answer.

The distribution computed by $\mathcal{B}$ and submitted to $\mathcal{A}$ equals either $([\![\varphi]\!]_{A_\eta})$ or $([\![\varphi']\!]_{A_\eta})$ depending on whichever $\mathcal{E}^*(\widehat{U},\widehat{a_0})$ is the encryption of $\widehat{U}$, or respectively, that of $\widehat{a_0}$ (in the latter case $\mathcal{E}^*(\widehat{U},\widehat{a_0}) = \mathcal{E}_{n_0}(\widehat{a_0})$ is simply a random number). Thus the probability that $\mathcal{B}$ guesses the right answer is the same as $\mathcal{A}$. Now it may happen that $\mathcal{B}$ does not meet the second requirement for winning the ω-IND-P1-C1 game, that is: (i) there exists a subterm $\mathsf{enc}_{n_0}(T,k)$ such that $T \neq U$ and $\widehat{T} \in \{\widehat{U},\widehat{a_0}\}$ or (ii) there exists a subterm $\mathsf{dec}_{n_0}(T,k)$ such that $\widehat{T} = \mathcal{E}^*(\widehat{U},\widehat{a_0})$.

For (i), the probability that $\widehat{T} = \widehat{a_0}$ is negligible by construction. Moreover, as T and $T^0 = \mathsf{enc}_{n_0}(U,k)$ are two subterms of φ and T^0 is not a subterm of T, the frame $\varphi' = \{x = T, y = U\}$ has no encryption cycles and $|\varphi'|_e < |\varphi|_e = n+1$. The induction hypothesis Q_n implies that the probability for $\widehat{T} = \widehat{U}$ is negligible.

As for (ii), if the challenge ciphertext $\mathcal{E}^*(\widehat{U},\widehat{a_0})$ is the encryption of its second argument, that is $\mathcal{E}_{n_0}(\widehat{a_0})$, then the probability for $\widehat{T} = \mathcal{E}^*(\widehat{U},\widehat{a_0})$ is negligible; otherwise $\mathcal{E}^*(\widehat{U},\widehat{a_0}) = \mathcal{E}_{n_0}(\widehat{U})$. Recall that $T^0 = \mathsf{enc}_{n_0}(U,k)$ is in $\mathcal{R}$-normal form, thus $U \neq \mathsf{dec}_{n_0}(T,k)$. As T^0 and $\mathsf{dec}_{n_0}(T,k)$ are two subterms of φ and T^0 is not a subterm of $\mathsf{dec}_{n_0}(T,k)$, the frame $\varphi' = \{x = U, y = \mathsf{dec}_{n_0}(T,k)\}$ has no encryption cycles and $|\varphi'|_e < |\varphi|_e = n+1$, hence the induction hypothesis Q_n implies that the probability for $\widehat{T} = \mathcal{E}_{n_0}(\widehat{U})$ is negligible.

To simplify the case analysis of (S3) and (S4), it is convenient to introduce the following lemma:

Lemma 13. *Let T_1, T_2 be two terms of sort $List_j$. Define for each $1 \leq i \leq j$, the i-th projection of a term T of sort $List_j$, by:*

$$\pi_i(T) = \mathsf{head}(\underbrace{\mathsf{tail}(\ldots\mathsf{tail}}_{i-1 \text{ times}}(T)))$$

Then (i) $T_1 =_{E_{\mathsf{sym}}} T_2$ *iff for all* $1 \leq i \leq j$, $\pi_i(T_1) =_{E_{\mathsf{sym}}} \pi_i(T_2)$ *and moreover (ii)* $\mathbb{P}\left[e1, e2 \leftarrow [\![T_1, T_2]\!]_{A_\eta}; e_1 = e_2\right]$ *is negligible iff for all* $1 \leq i \leq j$,

$$\mathbb{P}\left[e_1^i, e_2^i \leftarrow [\![\pi_i(T_1)\downarrow_{\mathcal{R}}, \pi_i(T_2)\downarrow_{\mathcal{R}}]\!]_{A_\eta}; e_1^i = e_2^i\right]$$

is negligible.

(The notation $T\downarrow_{\mathcal{R}}$ *stands for the* $\mathcal{R}$*-normal form of* T*.)*

Thanks to this lemma, it is sufficient to prove (S3) and (S4) for T_1 and T_2 of sort *Data* and in $\mathcal{R}$-normal form. (Indeed notice that if $\varphi = \{x = T_1, y = T_2\}$ has no encryption cycles, then $\varphi' = \{x' = \pi_i(T_1)\downarrow_{\mathcal{R}}, y' = \pi_i(T_2)\downarrow_{\mathcal{R}}\}$ has no encryption cycles and $|\varphi'|_e \leq |\varphi|_e$.)

Given the sorting system and the rewriting rules, a $\mathcal{R}$-reduced term T of sort *Data* may only be of the following forms:

1. a constant: 0 or 1,
2. a name of sort *Data*: $T = a$,
3. a projection of name of sort $List_j$: $T = \pi_i(a)$ $(1 \leq i \leq j)$,
4. a projection of a encryption/decryption of sort $List_j$: $T = \pi_i(\mathsf{enc}(U, V))$ with $U \notin \{\mathsf{dec}(T', V)\}$ or $T = \pi_i(\mathsf{dec}(U, V))$ with $U \notin \{\mathsf{enc}(T', V)\}$.

(S3) Q_0. As T_1 and T_2 contain no encryption/decryption symbol, only the cases 1–3 of the case analysis above can occur; the property follows directly.

(S4) $Q_{n+1} \Leftarrow (P_{n+1}$ and $Q_n)$. Let T_1 and T_2 be two distinct closed normal terms and $\varphi = \{x = T_1, y = T_2\}$. Assume that φ has no encryption cycles nor composed keys, and $|\varphi|_e = n + 1$.

1. If one of the two terms—say T_1— is of the form 1 (constant), 2 (name) or 3 (projection of a name). Then T_2 is of the form 4, for instance $T_2 = \pi_i(\mathsf{enc}(U, k))$ with $U \notin \{\mathsf{dec}(T', k)\}$.
   (a) If $T_1 \neq k$, by P_{n+1}, we have $([\![\varphi]\!]_{A_\eta}) \approx ([\![\{x = T_1, y = \pi_i(a)\}]\!]_{A_\eta})$ for some fresh name a. In particular, the probability for the two components x and y to be equal is negligible.
   (b) If $T_1 = k$, assume that T_1 and T_2 yields the same concrete value with significant probability. Let $List_{n_0}$ be the sort of U. We build an adversary $\mathcal{A}$ to the ω-IND-P1-C1 game as follows:
      i. for each name a of sort s appearing in T_2, draw a value $\widehat{a} \xleftarrow{R} [\![s]\!]_{A_\eta}$;
      ii. draw a value $\widehat{a_0} \xleftarrow{R} [\![s]\!]_{A_\eta}$ for some fresh name a_0 of sort $List_{n_0}$;
      iii. compute $\widehat{T_2}$ recursively as follows:

$$\begin{aligned}
\widehat{\mathsf{enc}_n(T, k)} &= \mathcal{E}_n(\widehat{T}) \text{ if } T \neq U \\
\widehat{\mathsf{enc}_{n_0}(U, k)} &= \mathcal{E}^*(\widehat{U}, \widehat{a_0}) \\
\widehat{\mathsf{dec}_n(T, k)} &= \mathcal{D}_n(\widehat{T}) \\
\widehat{f(V_1, \ldots, V_n)} &= [\![f]\!]_{A_\eta}(\widehat{V_1}, \ldots, \widehat{V_n}) \quad \text{in the remaining cases}
\end{aligned}$$

using the same conventions as before;

iv. if $\mathcal{E}_{n_0}(\widehat{U}, \widehat{T_2}) = \mathcal{E}^*(\widehat{U}, \widehat{a_0}))$, return 0, otherwise return 1.

$\mathcal{A}$ guesses the correct answer with non-negligible probability. As before, we use the property Q_n to conclude that its advantage is non-negligible.

2. Suppose $T_1 = \pi_{i_1}(\mathsf{enc}(u_1, k_1))$ and $T_2 = \pi_{i_2}(\mathsf{enc}(u_2, k_2))$ (the 3 other cases with decryption symbols are similar). As φ has no encryption cycle, we may assume for instance that k_1 is maximal for $<_\varphi$. Let T be a maximal subterm of the form $\mathsf{enc}(U, k_1)$ or $\mathsf{dec}(U, k_1)$ in φ. By P_{n+1}, we have $(\llbracket\varphi\rrbracket_{A_\eta}) \approx (\llbracket\varphi'\rrbracket_{A_\eta})$ where $\varphi' = \varphi\{T \mapsto a\} = \{x = T_1', y = T_2'\}$ for some fresh name a. We then apply Q_n to T_1' and T_2'. □

Proof (of Lemma 13). Point (i) is easily shown by induction on i, using the equations of E_{sym}. For (ii), notice that:

$$\mathbb{P}\left[e1, e2 \leftarrow \llbracket T_1, T_2\rrbracket_{A_\eta}; e_1 = e_2\right] \leq \sum_{i=1}^{j} \mathbb{P}\left[e_1^i, e_2^i \leftarrow \llbracket \pi_i(T_1), \pi_i(T_2)\rrbracket_{A_\eta}; e_1^i = e_2^i\right]$$

and

$$\forall i, \quad \mathbb{P}\left[e1, e2 \leftarrow \llbracket T_1, T_2\rrbracket_{A_\eta}; e_1 = e_2\right] \geq \mathbb{P}\left[e_1^i, e_2^i \leftarrow \llbracket \pi_i(T_1), \pi_i(T_2)\rrbracket_{A_\eta}; e_1^i = e_2^i\right]$$

Besides it is clear from the unconditional $=_{E_{\mathsf{sym}}}$-soundness, that for any T_1, T_2:

$$\mathbb{P}\left[e1, e2 \leftarrow \llbracket T_1, T_2\rrbracket_{A_\eta}; e_1 = e_2\right] = \mathbb{P}\left[e1, e2 \leftarrow \llbracket T_1 \downarrow_{\mathcal{R}}, T_2 \downarrow_{\mathcal{R}}\rrbracket_{A_\eta}; e_1 = e_2\right]$$

□

Proof (of Theorem 11). Thanks to the (unconditional) $=_{E_{\mathsf{sym}}}$-soundness, it is enough to prove the property on frames in $\mathcal{R}$-normal form.

We begin by proving the following lemma:

Lemma 14. *Assume that the concrete implementations for the encryption and its inverse satisfy both the ω-IND-P1-C1 assumption. For every well-formed $\mathcal{R}$-normal frame φ, $(\llbracket\varphi\rrbracket_{A_\eta}) \approx (\llbracket\overline{\varphi}\rrbracket_{A_\eta})$ where $\overline{\varphi}$ is the transparent frame associated to φ following the algorithmic proof of Proposition 9 (this transparent frame is uniquely defined modulo renaming of names.).*

Now recall that by Proposition 5 and since $\varphi \approx \overline{\varphi}$, we have:

$$\llbracket\overline{\varphi}\rrbracket_{A_\eta} = \llbracket\overline{\varphi}\rrbracket_{A_\eta}^{\text{ideal}} = \llbracket\varphi\rrbracket_{A_\eta}^{\text{ideal}}$$

Therefore the soundness criterion holds for well-formed $\mathcal{R}$-normal frames and we conclude by Proposition 3. □

Notice that the use of the ideal semantics could not be easily avoided as two statically equivalent transparent frames may not be equal modulo renaming of bound names.

Proof (of Lemma 14). We prove the property by induction on the number m of encryptions and decryptions by non-deducible keys in φ.

If $m = 0$, by the well-formedness condition, φ is already a transparent frame.

Suppose that $m > 0$. As φ has no encryption cycle, we choose a non-deducible (atomic) key k appearing in φ, such that k is maximal for the encryption relation $>_\varphi$.

As k is not deducible, is maximal for $>_\varphi$ and φ contains no head and tail symbols, the only occurrences of k in φ are as encryption or decryption keys. Let T be a maximal subterm of φ of the form $T = \mathsf{enc}(U, k)$ or $T = \mathsf{dec}(U, k)$. We apply Lemma 12 on φ and T and conclude by induction hypothesis on the obtained frame φ'. □

Note on the cryptographic assumptions.. Cryptographic assumptions of Theorem 11 may appear strong compared to existing work on passive adversaries [4, 13]. This seems unavoidable when we allow frames to contain both encryption and decryption symbols.

In the case where the two frames to be compared contain no decryption symbols, our proofs are easily adapted to work when the encryption scheme is ω-IND-P1-C0 only, where ω-IND-P1-C0 is defined similarly to ω-IND-P1-C1 except that the adversary has no access to the decryption oracle. Such an assumption is realizable in practice using a variable-input-length cipher [39, 38].

Finally, it should be possible to recover the classical assumption IND-P1-C1 by modeling the ECB mode (Electronic Code Book). Consider two new symbols $\mathsf{enc} : \mathit{Data} \times \mathit{Data} \to \mathit{Data}$ and $\mathsf{dec} : \mathit{Data} \times \mathit{Data} \to \mathit{Data}$, and define the symbols enc_n and dec_n (formally and concretely) recursively by

$$\begin{aligned} \mathsf{enc}_{n+1}(x, y) &= \mathsf{cons}_n(\mathsf{enc}(\mathsf{head}_n(x), y), \mathsf{enc}_n(\mathsf{tail}_n(x), y)) \quad \text{and} \\ \mathsf{dec}_{n+1}(x, y) &= \mathsf{cons}_n(\mathsf{dec}(\mathsf{head}_n(x), y), \mathsf{dec}_n(\mathsf{tail}_n(x), y)) \end{aligned}$$

together with the equations

$$\begin{aligned} \mathsf{dec}(\mathsf{enc}(x, y), y) &= x \\ \mathsf{enc}(\mathsf{dec}(x, y), y) &= y \end{aligned}$$

Define well-formed frames as those of which the normal forms contain no encryption cycles. Then, similar techniques can be applied to show that $\approx_{E_{\mathsf{sym}}}$-soundness holds for well-formed frames as soon as the implementations for enc and dec are both IND-P1-C1, or equivalently [37], enc is SPRP.

Note on the well-formedness assumptions.. We may also note that it is possible to slightly relax the assumptions of well-formedness of frames. In particular we could allow encryption cycles on deducible keys and for instance allow the frame $\{x = \mathsf{enc}(k_1, k_2),\ y = \mathsf{enc}(k_2, k_1),\ z = k_1\}$ which is currently discarded. As these extensions are not essential for our results we prefer to avoid unnecessary clutter and keep the definitions simple.

6. Conclusion and future work

In this paper we developed a general framework for relating formal and computational models of security protocols in the presence of a passive attacker. These are the first results on abstract models allowing arbitrary equational theories. We define the soundness and faithfulness of cryptographic implementations with respect to abstract models. We also provide a soundness criterion which is not only sufficient but also necessary for many theories. Finally, we provide new soundness results for the exclusive OR and a theory of ciphers and lists.

A direction for further work is to study the soundness of other theories. An interesting case is the combination of the two theories considered in this paper, that is modeling the exclusive OR, ciphers and lists. Another interesting open problem is to generalize the notion of transparent frames so as to include probabilistic encryption, while retaining the essential properties of transparent frames. Finally, an ambitious extension is to consider the case of an active attacker in presence of general equational theories.

Acknowledgments.. We would like to thank the anonymous reviewers for their helpful suggestions. This work was partially supported by the ACI JC 9005, the ARA SSIA FormaCrypt and the ANR SESUR AVOTÉ.

References

[1] M. Baudet, V. Cortier, S. Kremer, Computationally sound implementations of equational theories against passive adversaries, in: Proc. 32nd International Colloquium on Automata, Languages and Programming (ICALP'05), Vol. 3580 of LNCS, Springer, 2005, pp. 652–663.

[2] D. Dolev, A. C. Yao, On the security of public key protocols, IEEE Transactions on Information Theory IT-29 (12) (1983) 198–208.

[3] S. Goldwasser, S. Micali, Probabilistic encryption, Journal of Computer and System Sciences 28 (1984) 270–299.

[4] M. Abadi, P. Rogaway, Reconciling two views of cryptography (the computational soundness of formal encryption), in: Proc. 1st IFIP International Conference on Theoretical Computer Science (IFIP–TCS'00), Vol. 1872 of LNCS, 2000, pp. 3–22.

[5] H. Comon, V. Shmatikov, Is it possible to decide whether a cryptographic protocol is secure or not?, Journal of Telecommunications and Information Technology (4/2002) 5–15.

[6] V. Cortier, S. Delaune, P. Lafourcade, A survey of algebraic properties used in cryptographic protocols, Journal of Computer Security 14 (1) (2006) 1–43.

[7] M. Abadi, C. Fournet, Mobile values, new names, and secure communications, in: Proc. 28th Annual ACM Symposium on Principles of Programming Languages (POPL'01), 2001, pp. 104–115.

[8] R. Corin, J. Doumen, S. Etalle, Analysing password protocol security against off-line dictionary attacks, in: Proc. 2nd International Workshop on Security Issues with Petri Nets and other Computational Models (WISP'04), Vol. 121 of ENTCS, 2005, pp. 47–63.

[9] M. Baudet, Deciding security of protocols against off-line guessing attacks, in: Proc. 12th ACM Conference on Computer and Communications Security (CCS'05), ACM Press, 2005, pp. 16–25.

[10] M. Abadi, V. Cortier, Deciding knowledge in security protocols under equational theories, in: Proc. 31st International Colloquium on Automata, Languages and Programming (ICALP'04), Vol. 3142 of LNCS, 2004, pp. 46–58.

[11] B. Blanchet, Automatic proof of strong secrecy for security protocols, in: Proc. 25th IEEE Symposium on Security and Privacy (SSP'04), 2004, pp. 86–100.

[12] P. Adão, G. Bana, A. Scedrov, Computational and information-theoretic soundness and completeness of formal encryption, in: Proc. 18th IEEE Computer Security Foundations Workshop (CSFW'05), 2005, pp. 170–184.

[13] D. Micciancio, B. Warinschi, Completeness theorems for the Abadi-Rogaway logic of encrypted expressions, Journal of Computer Security 12 (1) (2004) 99–129.

[14] P. Laud, Computationally secure information flow, Ph.D. thesis, Universität des Saarlandes (2002).

[15] P. Laud, R. Corin, Sound computational interpretation of formal encryption with composed keys, in: Proc. 6th International Conference on Information Security and Cryptology (ICISC'03), Vol. 2971 of LNCS, 2004, pp. 55–66.

[16] P. Adão, J. Herzog, G. Bana, A. Scedrov, Soundness of formal encryption in the presence of key-cycles, in: Proc. 10th European Symposium on Research in Computer Security (ESORICS'05), Vol. 3679 of LNCS, 2005, pp. 374–396.

[17] M. Backes, B. Pfitzmann, M. Waidner, A composable cryptographic library with nested operations, in: Proc. 10th ACM Conference on Computer and Communications Security (CCS'03), ACM Press, 2003, pp. 220–230.

[18] M. Backes, B. Pfitzmann, Symmetric encryption in a simulatable Dolev-Yao style cryptographic library, in: Proc. 17th IEEE Computer Science Foundations Workshop (CSFW'04), 2004, pp. 204–218.

[19] M. Backes, B. Pfitzmann, M. Waidner, Symmetric authentication within simulatable cryptographic library, in: Proc. 8th European Symposium on Research in Computer Security (ESORICS'03), LNCS, 2003, pp. 271–290.

[20] V. Cortier, B. Warinschi, Computationally sound, automated proofs for security protocols, in: Proc. 14th European Symposium on Programming (ESOP'05), Vol. 3444 of LNCS, 2005, pp. 157–171.

[21] R. Janvier, Y. Lakhnech, L. Mazaré, Completing the picture: Soundness of formal encryption in the presence of active adversaries, in: Proc. 14th European Symposium on Programming (ESOP'05), Vol. 3444 of LNCS, 2005, pp. 172–185.

[22] R. Canetti, J. Herzog, Universally composable symbolic analysis of mutual authentication and key-exchange protocols (extended abstract), in: Proc. 3rd Theory of Cryptography Conference (TCC'06), Vol. 3876 of LNCS, 2006, pp. 380–403.

[23] P. Laud, Symmetric encryption in automatic analyses for confidentiality against active adversaries, in: Proc. IEEE Symposium on Security and Privacy (SSP'04), 2004, pp. 71–85.

[24] A. Datta, A. Derek, J. C. Mitchell, V. Shmatikov, M. Turuani, Probabilistic Polynomial-time Semantics for a Protocol Security Logic, in: Proc. 32nd International Colloquium on Automata, Languages and Programming, ICALP, Vol. 3580 of LNCS, Springer, 2005, pp. 16–29, lisboa, Portugal.

[25] B. Blanchet, A computationally sound mechanized prover for security protocols, in: IEEE Symposium on Security and Privacy, IEEE Computer Society Press, 2006, pp. 140–154.

[26] M. Backes, B. Pfitzmann, Limits of the cryptographic realization of dolev-yao-style xor, in: Proc. 10th European Symposium on Research in Computer Security (ESORICS'05), Vol. 3679 of LNCS, 2005, pp. 336–354.

[27] M. Abadi, B. Warinschi, Password-based encryption analyzed, in: Proc. 32nd International Colloquium on Automata, Languages and Programming (ICALP'05), Vol. 3580 of LNCS, 2005, pp. 664–676.

[28] M. Abadi, M. Baudet, B. Warinschi, Guessing attacks and the computational soundness of static equivalence, in: Proc. 9th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'06), Vol. 3921 of LNCS, 2006, pp. 398–412.

[29] G. Bana, P. Mohassel, T. Stegers, The computational soundness of formal indistinguishability and static equivalence, in: Proc. 11th Asian Computing Science Conference (ASIAN'06), Vol. 4435 of LNCS, Springer, 2006, pp. 182–196.

[30] S. Kremer, L. Mazaré, Adaptive soundness of static equivalence, in: Proc. 12th European Symposium on Research in Computer Security (ESORICS'07), Vol. 4734 of LNCS, Springer, 2007, pp. 610–625.

[31] H. Comon-Lundh, V. Cortier, Computational soundness of observational equivalence, in: Proc. 15th ACM Conference on Computer and Communications Security (CCS'08), ACM Press, 2008, pp. 109–118.

[32] M. Abadi, C. Fournet, Mobile values, new names, and secure communication, in: Proc. of the 28th ACM Symposium on Principles of Programming Languages (POPL'01), 2001, pp. 104–115.

[33] S. Hohenberger, The cryptographic impact of groups with infeasible inversion, Master's thesis, MIT (2003).

[34] R. L. Rivest, On the notion of pseudo-free groups, in: Proc. 1st Theory of Cryptography Conference (TCC'04), Vol. 2951 of LNCS, 2004, pp. 505–521.

[35] S. Goldwasser, M. Bellare, Lecture notes on cryptography (2008).

[36] D. Micciancio, The RSA group is pseudo-free, in: Advances in Cryptology – Proc. EUROCRYPT '05, Vol. 3494 of LNCS, 2005, pp. 387–403.

[37] D. H. Phan, D. Pointcheval, About the security of ciphers (semantic security and pseudo-random permutations), in: Proc. Selected Areas in Cryptography (SAC'04), Vol. 3357 of LNCS, 2004, pp. 185–200.

[38] S. Halevi, Invertible universal hashing and the TET encryption mode, in: Advances in Cryptology – Proc. CRYPTO '2007, Vol. 4622 of LNCS, 2007, pp. 412–429.

[39] M. Bellare, P. Rogaway, On the construction of variable-input-length ciphers, in: Proc. 6th Workshop on Fast Software Encryption (FSE'99), Vol. 1636 of LNCS, 1999, pp. 231–244.

[40] E. Contejean, C. Marché, B. Monate, X. Urbain, The CiME Rewrite Tool, `http://cime.lri.fr` (2000).

A. General results on static equivalence

We prove here some general properties of static equivalence concerning free symbols. We first establish a useful interpolation lemma.

Given a term $U = f(U_1, \ldots, U_n)$ where f is a free symbol (see Section 2.1) and a name a of the same sort as U, the *cutting function* $\mathsf{cut}_{U,a}$ is defined recursively as follows: $\mathsf{cut}_{U,a}(u) = u$ if u is a variable or a name, and

$$\mathsf{cut}_{U,a}(g(T_1, \ldots, T_k)) = \begin{cases} a & \text{if } g = f,\ k = n \text{ and } \forall 1 \leq i \leq n,\ U_i =_E T_i \\ g(\mathsf{cut}_{U,a}(T_1), \ldots, \mathsf{cut}_{U,a}(T_k)) & \text{otherwise} \end{cases}$$

Thus, the effect of function $\mathsf{cut}_{U,a}(T)$ is to substitute some (but not all) subterms of T equal to U modulo E with a.

Lemma 15. *Let $U = f(U_1, \ldots, U_n)$ be a term such that f is a free symbol. Let a be a name of the same sort as U. For any two terms M and N,*

$$M =_E N \quad \textit{implies} \quad \mathsf{cut}_{U,a}(M) =_E \mathsf{cut}_{U,a}(N).$$

Proof. By Birkhoff's theorem, $M =_E N$ means that there exist $n \geq 0$ and $M_0, \ldots, M_n$ such that $M = M_0 \leftrightarrow_E M_1 \leftrightarrow_E M_n = N$ where $\leftrightarrow_E$ denotes one step of rewriting along one equation in (the generating set of) E, oriented in either direction.

To prove the property by induction on n, it suffices to consider the case $n = 1$. More precisely, assume that there exists an equation $l = r$ in E, a position p and a substitution θ such that $M|_p = l\theta$ and $N = M[r\theta]_p$. By definition of free symbols, we may assume that f does not occur in l and r. We consider two cases depending on whether the cutting function $\mathsf{cut}_{U,a}$ cuts a subterm above p or not.

- Either there exists a proper prefix p' of p such that $M|_{p'} = f(T_1, \ldots, T_k)$ and for all i, $U_i =_E T_i$. We consider the smallest p' that satisfies this property. Thus $p = p' \cdot i \cdot p''$ and $N = M[f(T_1, \ldots, T_i[r\theta]_{p''}, \ldots, T_n)_{p'}]$. Both terms $f(T_1, \ldots, T_k)$ and $f(T_1, \ldots, T_i[r\theta]_{p''}, \ldots, T_n)$ are substituted with a, thus $\mathsf{cut}_{U,a}(M) = \mathsf{cut}_{U,a}(N)$.
- Or no such cutting position p' is a proper prefix of p. This means that $\mathsf{cut}_{U,a}(M[x]_p) = \mathsf{cut}_{U,a}(N[x]_p)$ and $\mathsf{cut}_{U,a}(M) = \mathsf{cut}_{U,a}(M[x]_p)[\mathsf{cut}_{U,a}(l\theta)]_p$, where x is a fresh variable. Moreover, $\mathsf{cut}_{U,a}(l\theta) = l\mathsf{cut}_{U,a}(\theta)$ and $\mathsf{cut}_{U,a}(r\theta) = r\mathsf{cut}_{U,a}(\theta)$ since f is free. We deduce

$$\begin{array}{rcl} \mathsf{cut}_{U,a}(M) & = & \mathsf{cut}_{U,a}(M[x]_p)[\mathsf{cut}_{U,a}(l\theta)]_p \\ & = & \mathsf{cut}_{U,a}(N[x]_p)[l\mathsf{cut}_{U,a}(\theta)]_p \\ & =_E & \mathsf{cut}_{U,a}(N[x]_p)[r\mathsf{cut}_{U,a}(\theta)]_p \\ & = & \mathsf{cut}_{U,a}(N) \end{array}$$

Using this lemma, we establish two simple properties of free symbols.

Corollary 16. *Let f be a free symbol and $f(T_1, \ldots, T_n)$ a term of a non-degenerated type τ.*

1. *For every $U_1, \ldots, U_n$ of the appropriate sort,*

$$f(T_1, \ldots T_n) =_E f(U_1, \ldots, U_n) \quad \textit{iff} \quad \forall i,\ T_i =_E U_i.$$

2. *Let U be a term of sort τ such that f does not appear in U. Then*

$$f(T_1, \ldots, T_n) \neq_E U.$$

Proof.

1. The right-to-left implication is trivial. Let $T = f(T_1, \dots T_n)$ and $U = f(U_1, \dots, U_n)$. By contradiction, assume that there exists an i such that $T_i \neq_E U_i$. Let a_1, a_2 be two fresh names of sort τ. We apply Lemma 15 on the equation $T =_E U$ successively with cut_{T,a_1} and cut_{U',a_2} where $U' = \mathsf{cut}_{T,a_1}(U) = f(\mathsf{cut}_{T,a_1}(U_1), \dots \mathsf{cut}_{T,a_1}(U_n))$. We obtain $a_1 =_E a_2$, hence τ is degenerated; contradiction.

2. Assume $f(T_1, \dots, T_n) =_E U$. Then by Lemma 15, since f does not occur in U, we obtain $a =_E U$ for some fresh name a, hence τ is degenerated; contradiction.

We are now ready to prove our propositions.

Proposition 17. *Let T_1, T_2 be two terms of sort s such that $T_1 \neq_E T_2$. Assume a free symbol $\mathsf{h}_s : s \times Key \to Hash$ such that the sort Key is not degenerated. Consider the frame $\varphi_1 = \{x_1 = \mathsf{h}_s(T_1, k),\ x_2 = \mathsf{h}_s(T_2, k)\}$ where k is a fresh name. Let $\varphi_2 = \{x_1 = n, x_2 = n'\}$ where n, n' are two distinct fresh names of sort Hash. Then we have $\varphi_1 \approx_E \varphi_2$.*

Proof. Let M and N be two terms such that $\mathrm{var}(M, N) \subseteq \mathrm{dom}(\varphi)$ and $\mathrm{names}(M, N) \cap \mathrm{names}(\varphi) = \emptyset$.

Assume $M\varphi_2 =_E N\varphi_2$. Let θ be the substitution $\{n \mapsto \mathsf{h}_s(T_1, k), n' \mapsto \mathsf{h}_s(T_2, k)\}$. Since the equational theory E is stable by substitution of names, we have $M\varphi_2\theta =_E N\varphi_2\theta$, hence, $M\varphi_1 =_E N\varphi_1$ as n, n' are fresh names.

Conversely, assume $M\varphi_1 =_E N\varphi_1$. Let $U_1 = \mathsf{h}_s(T_1, k)$. By Lemma 15, we have $\mathsf{cut}_{U_1,n}(M\varphi_1) =_E \mathsf{cut}_{U_1,n}(N\varphi_1)$. Since k does not appear in M nor N, by Corollary 16, it holds that $\mathsf{cut}_{U_1,n}(M\varphi_1) = M\mathsf{cut}_{U_1,n}(\varphi_1)$ and $\mathsf{cut}_{U_1,n}(N\varphi_1) = N\mathsf{cut}_{U_1,n}(\varphi_1)$. Now, using $T_1 \neq_E T_2$, we prove $\mathsf{cut}_{U_1,n}(\varphi_1) = \{x_1 = n,\ x_2 = \mathsf{h}_s(T_2, k)\}$. Indeed, we have $\mathsf{cut}_{U_1,n}(\mathsf{h}_s(T_2, k)) = \mathsf{h}_s(\mathsf{cut}_{U_1,n}(T_2), k)$ since $T_1 \neq_E T_2$. Besides, as k does not appear in T_2, by Corollary 16, we have $\mathsf{cut}_{U_1,n}(T_2) = T_2$. Similarly, by applying $\mathsf{cut}_{U_2,n'}$ with $U_2 = \mathsf{h}_s(T_2, k)$, we obtain

$$M\mathsf{cut}_{U_2,n'}(\mathsf{cut}_{U_1,n}(\varphi_1)) =_E N\mathsf{cut}_{U_2,n'}(\mathsf{cut}_{U_1,n}(\varphi_1)),$$

that is, $M\varphi_2 =_E N\varphi_2$. □

Proposition 18. *Let φ be a frame and T a term of sort s. Assume a free symbol $\mathsf{h}_s : s \times Key \to Hash$ such that the sort Key is not degenerated. Let $\varphi_1 = \varphi \cup \{x = \mathsf{h}_s(T, k), y = k\}$ and $\varphi_2 = \varphi \cup \{x = n, y = k\}$ where x, y are fresh variables, k is a fresh name of sort Key, n is a fresh name of sort Hash. If $\varphi \not\vdash_E T$, then $\varphi_1 \approx_E \varphi_2$.*

Proof. Let M and N be two terms such that $\mathrm{var}(M, N) \subseteq \mathrm{dom}(\varphi)$ and $\mathrm{names}(M, N) \cap \mathrm{names}(\varphi) = \emptyset$. We prove that $M\varphi_2 =_E N\varphi_2$ implies $M\varphi_1 =_E N\varphi_1$ similarly as for Proposition 17.

Conversely, assume $M\varphi_1 =_E N\varphi_1$. Let $U = \mathsf{h}_s(T, k)$. By Lemma 15, we have $\mathsf{cut}_{U,n}(M\varphi_1) =_E \mathsf{cut}_{U,n}(N\varphi_1)$. Let us prove that $\mathsf{cut}_{U,n}(M\varphi_1) = M\mathsf{cut}_{U,n}(\varphi_1)$.

Indeed, otherwise, there exists a subterm M_1 of M such that M_1 is not a variable and $M_1\varphi_1 = \mathsf{h}_s(T', T'')$ with $T' =_E T$ and $T'' =_E k$. Since M_1 is not a variable, M_1 is of the form $M_1 = \mathsf{h}_s(M_1', M_1'')$ with $M_1'\varphi_1 = T' =_E T$, which implies that T is deducible; contradiction.

We deduce that $\mathsf{cut}_{U,n}(M\varphi_1) = M\mathsf{cut}_{U,n}(\varphi_1)$, and similarly $\mathsf{cut}_{U,n}(N\varphi_1) = N\mathsf{cut}_{U,n}(\varphi_1)$. Thus $M\mathsf{cut}_{U,n}(\varphi_1) =_E N\mathsf{cut}_{U,n}(\varphi_1)$. By Corollary 16, as k does not appear in φ, we have that $\mathsf{cut}_{U,n}(\varphi) = \varphi$, hence $\mathsf{cut}_{U,n}(\varphi_1) = \varphi_2$ and $M\varphi_2 =_E N\varphi_2$. □

Proposition 19. *Let T_1, T_2 be two terms of sort s such that $T_1 =_E T_2$. Assume a free symbol $\mathsf{h}_s : s \times Key \to Hash$ such that Key is not degenerated. Let $\varphi = \{x_1 = \mathsf{h}_s(T_1, k),\ x_2 = \mathsf{h}_s(T_2, k)\}$. Then, $\varphi \approx_E \{x_1 = n, x_2 = n\}$ where n is a fresh name of sort Hash.*

Proof. Let M and N be two terms such that $\mathrm{var}(M, N) \subseteq \mathrm{dom}(\varphi)$ and $\mathrm{names}(M, N) \cap \mathrm{names}(\varphi) = \emptyset$. We prove that $M\varphi_2 =_E N\varphi_2$ implies $M\varphi_1 =_E N\varphi_1$ similarly as for Proposition 17.

Conversely, assume $M\varphi_1 =_E N\varphi_1$. Let $U = \mathsf{h}_s(T_1, k)$. By Lemma 15, we have $\mathsf{cut}_{U,n}(M\varphi_1) =_E \mathsf{cut}_{U,n}(N\varphi_1)$. Since k does not appear in M nor N, by Corollary 16, we have $\mathsf{cut}_{U,n}(M\varphi_1) = M\mathsf{cut}_{U,n}(\varphi_1)$ and $\mathsf{cut}_{U,n}(N\varphi_1) = N\mathsf{cut}_{U,n}(\varphi_1)$. Now, since $T_1 =_E T_2$, we obtain $\mathsf{cut}_{U,n}(\varphi_1) = \{x_1 = n,\ x_2 = n\} = \varphi_2$. Thus we have $M\varphi_2 =_E N\varphi_2$. □

B. Static equivalence in groups

We establish some properties of static equivalence in the equational theory of Abelian groups E_{G} defined in Section 3.2. For this purpose we characterize equivalence classes in E_{G} by a representation lemma.

Let $\mathcal{X}_A$ ($\mathcal{X}_G$ and $\mathcal{X}_{Hash}$ respectively) be the set of variables of sort A (G and *Hash* respectively). Let $\mathcal{N}_A$ ($\mathcal{N}_G$ and $\mathcal{N}_{Hash}$ respectively) be the set of names of sort A (G and *Hash* respectively). Let AC be the equational theory corresponding to the subset of equations from E_{G}, modeling the associativity and commutativity of the three operators $\cdot$, $+$ and $*$.

We call *unitary monomial of sort A* a function $\beta : \mathcal{X}_A \cup \mathcal{N}_A \to \mathbb{N}$ almost everywhere zero, *i*.e., except for a finite number of entries. Such a function β can be considered as a term of sort A (modulo AC):

$$\beta =_{AC} \prod_{a \in \mathcal{N}_A,\ \beta(a) \neq 0} a^{\beta(a)} \cdot \prod_{u \in \mathcal{X}_A,\ \beta(u) \neq 0} u^{\beta(u)}$$

where empty products are considered to be the term 1_A, and $a^{\beta(a)}$ ($\beta(a) \neq 0$) denotes the term $\underbrace{a \cdot \ldots \cdot a}_{\beta(a) \text{ times}}$. We denote $\mathcal{M}_A$ the set of all unitary monomials of sort A.

A *canonical form of sort* A is a function $\alpha : \mathcal{M}_A \to \mathbb{Z}$ almost everywhere zero. We consider such a function α as a term of sort A (modulo AC):

$$\alpha =_{AC} \sum_{\beta \in \mathcal{M}_A,\ \alpha(\beta) \neq 0} \alpha(\beta) \cdot \beta$$

where empty sums are considered to be the term 0_A, and integers are naturally represented as 0_A, $1_A + \ldots + 1_A$ or $-(1_A + \ldots + 1_A)$ of sort A.

A *canonical form of sort* G is a function γ, mapping terms in $\mathcal{X}_N \cup \mathcal{N}_N$ to canonical forms of sort A, almost everywhere zero, *i.e.*, the function evaluates to the constant 0 except for a finite number of entries. We consider a canonical form γ to be a term of sort G (modulo AC):

$$\gamma =_{AC} \prod_{g \in \mathcal{N}_G,\ \gamma(g) \neq 0} g^{\gamma(g)} \ * \prod_{x \in \mathcal{X}_G,\ \gamma(x) \neq 0} x^{\gamma(x)}$$

where empty products are considered to be equal to 1_G.

A *canonical form of sort Hash*, denoted ι, is either a variable of sort *Hash* : $\iota = z \in \mathcal{X}_{Hash}$, a name of sort *Hash* : $\iota = h \in \mathcal{N}_{Hash}$ or a canonical form γ of sort G considered to be a term $\iota = \mathsf{h}(\gamma)$.

Lemma 20. *For any term T of sort A (G, Hash, respectively), there exists a unique canonical form α_T (γ_T, ι_T, respectively) such that*

$$T =_{E_\mathsf{G}} \alpha_T$$

($T =_{E_\mathsf{G}} \gamma_T$, $T =_{E_\mathsf{G}} \iota_T$, respectively).

Proof (Sketch). We show the existence of a canonical form of a term T by induction on the structure of T. For instance, given $T = T_1 * T_2$, and two canonical forms α_{T_1} and α_{T_2}, we obtain the canonical form of T by rearranging the product $\alpha_{T_1} * \alpha_{T_2}$ modulo E_G (and if necessary the induction hypthesis is also used on the exponents). To show the uniqueness of the normal form, it is sufficient to show that whenever two canonical terms are equal as terms modulo E_G, they are also equal "mathematically". Formally this is established by studying the AC normal form of each canonical form with respect to the following AC-convergent rewriting system.

$$\begin{array}{rcl@{\qquad}rcl}
u + 0_A & \to & u & x * 1_G & \to & x \\
u + (-u) & \to & 0_A & (x^u)^v & \to & x^{(u \cdot v)} \\
u \cdot 1_A & \to & u & x^u * x^v & \to & x^{u+v} \\
(u+v) \cdot w & \to & u \cdot w + v \cdot w & x^{1_A} & \to & x \\
u \cdot 0_A & \to & 0_A & x^{0_A} & \to & 1_G \\
-(u+v) & \to & (-u) + (-v) & (x * y)^u & \to & x^u * y^u \\
(-u) \cdot v & \to & -(u \cdot v) & x * x & \to & x^{(1_A + 1_A)} \\
-(-u) & \to & u & x * x^u & \to & x^{u+1_A} \\
-0_A & \to & 0_A & (1_G)^u & \to & 1_G
\end{array}$$

This rewriting system has been obtained by orienting and completing the equations generating E_{G}, except AC, using the tool Cime [40]. □

Proposition 21. *Let* $\varphi_1 = \nu g, a, b.\{x_1 = g, x_2 = g^a, x_3 = g^b, x_4 = g^{a\cdot b}\}$ *and* $\varphi_2 = \nu g, a, b, c.\{x_1 = g, x_2 = g^a, x_3 = g^b, x_4 = g^c\}$. *We have that* $\varphi_1 \approx_{E_{\mathsf{G}}} \varphi_2$.

Proof. Let M, N be two terms of the same sort such that $\mathrm{var}(M, N) \subseteq \mathrm{dom}(\varphi_1)$ and $\mathrm{names}(M, N) \cap \mathrm{names}(\varphi_1, \varphi_2) = \emptyset$.

Assume $M\varphi_2 =_{E_{\mathsf{G}}} N\varphi_2$. Let θ be the substitution $\{c \mapsto a \cdot b\}$. Since the equational theory E is stable by substitution of names, we have $M\varphi_2\theta =_E N\varphi_2\theta$, that is, $M\varphi_1 =_{E_{\mathsf{G}}} N\varphi_1$ since $c \notin \mathrm{names}(M, N)$.

Conversely, assume $M\varphi_1 =_{E_{\mathsf{G}}} N\varphi_1$. If M and N are of sort A, then $\mathrm{var}(M, N) = \emptyset$ and hence $M\varphi_2 = M\varphi_1 =_E N\varphi_1 = N\varphi_2$.

Otherwise, M and N are of sort G. As $M\varphi_1 =_{E_{\mathsf{G}}} N\varphi_1$ is equivalent to $M\varphi_1 * (N\varphi_1)^{-1_A} = 1_G$, we suppose that $N = 1_G$.

As $\mathrm{var}(M) \subseteq \mathrm{dom}(\varphi_1)$ and $\mathrm{names}(M) \cap \mathrm{names}(\varphi_1, \varphi_2) = \emptyset$, the canonical form γ of M is of the form

$$M =_{E_{\mathsf{G}}} \prod_{g' \neq g} {g'}^{\gamma(g')} * x_1^{\gamma(x_1)} * \ldots * x_4^{\gamma(x_4)}$$

where $\gamma(g')$ and $\gamma(x_i)$ represent closed terms with disjoint names $\{a, b, c\}$. Hence, we have that

$$M\varphi_1 =_{E_{\mathsf{G}}} \prod_{g' \neq g} {g'}^{\gamma(g')} * g^{\gamma(x_1) + \gamma(x_2)\cdot a + \gamma(x_3)\cdot b + \gamma(x_4) a\cdot b} =_{E_{\mathsf{G}}} 1_G$$

and we conclude that for any i, $\gamma(x_i) = 0_A$ and for any g', $\gamma(g') = 0_A$, *i.e.*, $M = 1_G$. □

Proposition 22. *Let the frame* $\varphi_1 = \nu g, a.\{x_1 = g^a, x_2 = a, x_3 = \mathsf{h}(g)\}$ *and the frame* $\varphi_2 = \nu g, a, h.\{x_1 = g^a, x_2 = a, x_3 = h\}$. *We have that* $\varphi_1 \approx_{E_{\mathsf{G}}} \varphi_2$.

Proof. Let M, N be two terms such that $\mathrm{var}(M, N) \subseteq \mathrm{dom}(\varphi_1)$ and $\mathrm{names}(M, N) \cap \mathrm{names}(\varphi_1, \varphi_2) = \emptyset$.

Assume $M\varphi_2 =_E N\varphi_2$. Let θ be the substitution $\{h \mapsto \mathsf{h}(g)\}$. Since the equational theory E is stable by substitution of names, we have $M\varphi_2\theta =_E N\varphi_2\theta$, hence, as $h \notin \mathrm{names}(M, N)$, $M\varphi_1 =_E N\varphi_1$.

Conversely, assume that $M\varphi_1 =_{E_{\mathsf{G}}} N\varphi_1$. If M and N are of sort A or G, then $\mathrm{var}(M, N) \subseteq \{x_1, x_2\}$ and hence $M\varphi_2 = M\varphi_1 =_E N\varphi_1 = N\varphi_2$.

Otherwise, M and N are of sort *Hash*. We suppose that $M = x_3$ and $N = \mathsf{h}(N')$ where $\mathrm{var}(N') \subseteq \{x_1, x_2\}$ (other cases are trivial). As h is a free symbol, by Corollary 16, $M\varphi_1 =_{E_{\mathsf{G}}} N\varphi_1$ is equivalent to $N'\varphi_1 =_{E_{\mathsf{G}}} g$.

Given that $\mathrm{var}(N') \subseteq \{x_1, x_2\}$ and $\mathrm{names}(N') \cap \mathrm{names}(\varphi_1, \varphi_2) = \emptyset$, the canonical form γ of N' is of the form

$$N' =_{E_{\mathsf{G}}} \prod_{g' \neq g} {g'}^{\gamma(g')} * x_1^{\gamma(x_1)}$$

where $\gamma(g')$ and $\gamma(x_1)$ are terms that have no variable other than x_2 and do not contain a. Hence we have

$$N'\varphi_1 =_{E_{\mathsf{G}}} \prod_{g' \neq g} g'^{\gamma(g')\{x_2 \mapsto a\}} * g^{\gamma(x_1)\cdot a}$$

which contradicts $N'\varphi_1 =_{E_{\mathsf{G}}} g$. □

C. Static equivalence in ciphers and lists

Before proving Lemma 10, we first introduce a handy lemma to characterize deducible terms.

Lemma 23. *Let $\varphi = \nu\tilde{n}.\sigma$ be a closed frame in $\mathcal{R}$-normal form and T a term in $\mathcal{R}$-normal form. If $\varphi \vdash_{E_{\mathsf{sym}}} T$ then $T = C[T_1, \ldots, T_k]$ where the T_i are deducible subterms of φ and C is a context that does not contain private names that is* $\mathrm{names}(C) \cap \tilde{n} = \emptyset$.

Proof. By definition, $\varphi \vdash_{E_{\mathsf{sym}}} T$ if and only if there exists a term M such that $\mathrm{names}(M) \cap \mathrm{names}(\varphi) = \emptyset$ and $M\varphi =_{E_{\mathsf{sym}}} T$, that is, $M\varphi \to^*_{\mathcal{R}} T$. We prove Lemma 23 by induction on the size of M. The base case $M = x_i$ is trivial.

If $M = f(M_1, \ldots, M_k)$. We only consider the case where $M = \mathsf{dec}(M_1, M_2)$ since the other cases are similar. We have $M_1 \to^*_{\mathcal{R}} T_1$ and $M_2 \to^*_{\mathcal{R}} T_2$. By applying the induction hypothesis to M_1 and M_2, we obtain that $T_1 = C_1[T'_1, \ldots, T'_k]$ and $T_2 = C_2[T'_1, \ldots, T'_k]$ where the T'_i are deducible subterms of φ and C_1, C_2 are contexts that do not contain names. We have $M\varphi \to^*_{\mathcal{R}} \mathsf{dec}(T_1, T_2)$. Either $\mathsf{dec}(T_1, T_2)$ is in $\mathcal{R}$-normal form. In that case and by convergence of $\mathcal{R}$, we have $T = \mathsf{dec}(T_1, T_2)$, hence the result. Or $\mathsf{dec}(T_1, T_2)$ is not in $\mathcal{R}$-normal form. By convergence, we have $\mathsf{dec}(T_1, T_2) \to_{\mathcal{R}} T$. Since T_1 and T_2 are already in normal form, we must have $T_1 = \mathsf{enc}(T'_1, T_2)$ and $T = T'_1$. Either $C_1 = \mathsf{enc}(C'_1, C''_1)$ and we have $T = C'_1[T'_1, \ldots, T'_k]$. Or $C_1 = _$, which means that T_1 is a deducible subterm of φ. We deduce that T is a deducible subterm of φ, hence the result. □

We can now start the proof of Lemma 10.

Proof. In what follows, we say that a term or a context is *public* if it does not contain the names occurring in φ. Since $\varphi = \varphi'\{n \mapsto T\}$ and E_{sym} is stable by substitutions of names, we have $\mathrm{eq}_{E_{\mathsf{sym}}}(\varphi') \subseteq \mathrm{eq}_{E_{\mathsf{sym}}}(\varphi)$. To prove $\mathrm{eq}_{E_{\mathsf{sym}}}(\varphi) \subseteq \mathrm{eq}_{E_{\mathsf{sym}}}(\varphi')$, we introduce the following lemma. We set θ to be $\{n \mapsto T\}$. Let $n_1, \ldots, n_p$ be the names occurring in φ'.

Lemma 24. *Let C_1 be a context such that we have $\varphi' \vdash_{E_{\mathsf{sym}}} C_1[n_1, \ldots, n_p]$ and $C_1[n_1, \ldots, n_p]\theta \to_{\mathcal{R}} T$. Then there exists a public context C_2 such that $C_1 \to_{\mathcal{R}} C_2$ and $T = C_2[n_1, \ldots, n_p]\theta$.*

The lemma is proved by inspection of the rules of $\mathcal{R}$. The reduction occurs at some position p: the reduction $C_1[n_1,\ldots,n_p]|_p\theta \to_{\mathcal{R}} T$ occurs in head. Let $C_1'[n_1,\ldots,n_p] = C_1[n_1,\ldots,n_p]|_p$ If C_1' is itself an instance of the left-hand-side of a rule of $\mathcal{R}$, than we clearly have that $C_1' \to_{\mathcal{R}} C_2'$ such that $T = C_2[n_1,\ldots,n_p]\theta$, where C_2 is obtained from C_1 by replacing C_1' with C_2' at position p. If C_1' is not an instance of the left-hand-side of a rule of $\mathcal{R}$ and since T is already in $\mathcal{R}$-normal form, there are only four possibilities for $C_1'[n_1,\ldots,n_p]$.

- $C_1'[n_1,\ldots,n_p] = \mathsf{enc}(n_i, C_1''[n_1,\ldots,n_p])$. It must be the case that $n_i = n$, T is of the form $\mathsf{dec}(U,V)$ and $V = C_1''[n_1,\ldots,n_p]$. From Lemma 23 and since $\varphi' \vdash_{E_{\mathsf{sym}}} C_1[n_1,\ldots,n_p]$, either $C_1'[n_1,\ldots,n_p]$ is subterm of φ' or n_i and $C_1''[n_1,\ldots,n_p]$ are deducible. In both cases, we obtain a contradiction. Indeed, if $C_1'[n_1,\ldots,n_p]$ is subterm of φ' then $C_1'[n_1,\ldots,n_p]\theta = \mathsf{enc}(\mathsf{dec}(U,V),n_j)$ is a subterm of φ, which contradicts that φ is in normal form. If n_i and $C_1''[n_1,\ldots,n_p]$ are deducible then this contradicts $\varphi \not\vdash_{E_{\mathsf{sym}}} V$.

- $C_1'[n_1,\ldots,n_p] = \mathsf{dec}(n_i,n_j)$. This case is very similar to the previous one.

- $C_1'[n_1,\ldots,n_p] = \mathsf{cons}(n_i, C_1''[n_1,\ldots,n_p])$. It must be the case that $n_i = n$, T is of the form $\mathsf{head}(V)$ and $C_1''[n_1,\ldots,n_p] = \mathsf{tail}(V)$. From Lemma 23 and since $\varphi' \vdash_{E_{\mathsf{sym}}} C_1[n_1,\ldots,n_p]$, either $C_1'[n_1,\ldots,n_p]$ is subterm of φ' or n_i and $C_1''[n_1,\ldots,n_p]$ are deducible. As previously, in both cases, we obtain a contradiction. if $C_1'[n_1,\ldots,n_p]$ is subterm of φ' then $C_1'[n_1,\ldots,n_p]\theta = \mathsf{cons}(\mathsf{head}(V),\mathsf{tail}(V))$ is a subterm of φ, which contradicts that φ is in normal form. If n_i and $C_1''[n_1,\ldots,n_p]$ are deducible then both n and $\mathsf{tail}(V)$ are deducible in φ', which means that both $\mathsf{head}(V)$ and $\mathsf{tail}(v)$ are deducible in φ, thus V is deducible in φ, contradiction.

- $C_1'[n_1,\ldots,n_p] = \mathsf{cons}(C_1''[n_1,\ldots,n_p], n_i)$. This case is very similar to the previous one.

Now, let $(M = N) \in \mathrm{eq}_{E_{\mathsf{sym}}}(\varphi)$ and let us show that $(M = N) \in \mathrm{eq}_{E_{\mathsf{sym}}}(\varphi')$. We have $M\varphi =_{E_{\mathsf{sym}}} N\varphi$, that is, $M\varphi'\theta =_{E_{\mathsf{sym}}} N\varphi'\theta$. By convergence of $\mathcal{R}$, there exists a term T such that $M\varphi'\theta \to_{\mathcal{R}}^* T$ and $N\varphi'\theta \to_{\mathcal{R}}^* T$. By applying repeatedly Lemma 24, we obtain that $M\varphi' \to_{\mathcal{R}}^* T_1$ such that $T = T_1\theta$ and $N\varphi' \to_{\mathcal{R}}^* T_2$ such that $T = T_2\theta$. Assume that we have proved that $T_1 = T_2$. Then we have $M\varphi' =_{E_{\mathsf{sym}}} N\varphi'$, that is, $(M = N) \in \mathrm{eq}_{E_{\mathsf{sym}}}(\varphi')$, which concludes the proof. It remains for us to prove the following lemma.

Lemma 25. *Let T_1 and T_2 be two terms such that each T_i is either deducible from φ', that is, $\varphi' \vdash_{E_{\mathsf{sym}}} T_i$, or T_i is a subterm of φ'. Then $T_1\theta = T_2\theta$ implies $T_1 = T_2$.*

The lemma is proved by induction on the sum of the size of T_1 and T_2. First notice that, by Lemma 23, any subterm T' of one of the T_i verifies that T' is deducible from φ' or T' is a subterm of φ'.

- The base case is trivial.
- If none of T_1 or T_2 is n: $T_1 = f(T'_1, \ldots, T'_k)$ and $T_2 = f(T''_1, \ldots, T''_k)$. We must have $T'_i\theta = T''_i\theta$ for every $1 \leq i \leq k$. By applying the induction hypothesis, we obtain $T'_i = T''_i$ thus $T_1 = T_2$.
- The most difficult case is when $T_1 = n$ and $T_2 = f(T'_1, \ldots, T'_k)$. We first notice that since $n\theta = f(T'_1, \ldots, T'_k)\theta$, n cannot occur in T_2, thus $T_2 = T_2\theta$. Either T_2 is a subterm of φ', which is impossible by construction of φ' or T_2 deducible. Since T_2 is not a subterm of φ' and applying again Lemma 23, we get that the immediate subterms of T_2 are deducible in φ' (thus in φ), which contradicts the choice of T. □