Practical Everlasting Privacy

Myrto Arapinis¹, Véronique Cortier², Steve Kremer², and Mark Ryan¹

¹ School of Computer Science, University of Birmingham

² LORIA, CNRS, France

Abstract. Will my vote remain secret in 20 years? This is a natural question in the context of electronic voting, where encrypted votes may be published on a bulletin board for verifiability purposes, but the strength of the encryption is eroded with the passage of time. The question has been addressed through a property referred to as *everlasting privacy*. Perfect everlasting privacy may be difficult or even impossible to achieve, in particular in remote electronic elections. In this paper, we propose a definition of *practical everlasting privacy*. The key idea is that in the future, an attacker will be more powerful in terms of computation (he may be able to break the cryptography) but less powerful in terms of the data he can operate on (transactions between a vote client and the vote server may not have been stored).

We formalize our definition of everlasting privacy in the applied- π calculus. We provide the means to characterize what an attacker can break in the future in several cases. In particular, we model this for perfectly hiding and computationally binding primitives (or the converse), such as Pedersen commitments, and for symmetric and asymmetric encryption primitives. We adapt existing tools, in order to allow us to automatically prove everlasting privacy. As an illustration, we show that several variants of Helios (including Helios with Pedersen commitments) and a protocol by Moran and Naor achieve practical everlasting privacy, using the ProVerif and the AKiSs tools.

1 Introduction

Electronic voting schemes such as Helios [2], JCJ/Civitas [14,8], and Prêt-à-Voter [7] aim simultaneously to guarantee *vote privacy* (that is, the link between the voter and her vote will not be revealed), and *outcome verifiability* (that is, voters and observers can check that the declared outcome is correct). A common way to achieve verifiability is to publish a “bulletin board” that contains all encrypted votes (indeed, this is the way it is done in the systems cited above). The strength and key-length of the encryption should be chosen so that decryption by an attacker is impossible for as long as the votes are expected to remain private. To prevent coercion reprisal not just to the voter but also to her descendants, one may want vote privacy for up to 100 years.

Unfortunately, however, it is not possible to predict in any reliable way how long present-day encryptions will last. Weaknesses may be found in encryption algorithms, and computers will certainly continue to get faster. A coercer can plausibly assert that a voter should follow the coercer’s wishes because the bulletin board will reveal in (say) 10 years whether the voter followed the coercer’s instructions. For this reason, systems with “everlasting privacy” have been introduced by [18]. These systems do not rely

on encryptions whose strength may be eroded, but on commitments that are *perfectly* or *information-theoretically hiding*. These systems have computational verifiability instead of perfect verifiability, and are considered less usable and computationally more expensive than systems relying on encryptions. More recently, schemes have been proposed with a weaker form of everlasting privacy (e.g., [10,12]); they rely on encryptions for counting votes, but use commitments rather than encryptions for verifiability purposes. Thus, the bulletin board which only publishes the commitments does not weaken the privacy provided by the underlying scheme. Although the encrypted votes must be sent to the election administrators, it is assumed that these communications cannot be intercepted and stored *en masse*. We call this weaker form of everlasting privacy *practical everlasting privacy*.

Symbolic models for security protocol analysis have been used to model both privacy properties (e.g., [11,3,13]) and verifiability properties (e.g., [16,17]) of voting systems, but they are currently not capable of distinguishing *perfect* versus *computational* notions of privacy, or indeed, of verifiability. Our aim in this paper is to extend the model to allow these distinctions. We focus on practical everlasting privacy, and use our definitions to verify whether particular schemes satisfy that property.

Our contributions. Our first and main contribution is a general and practical definition of everlasting privacy. The key idea is that, in the future, an attacker will be more powerful in terms of computation (he may be able to break cryptography) but less powerful in terms of the data he can operate on (transactions between a vote client and the vote server may not have been stored). We therefore distinguish between standard communication channels (on which eavesdropping may be possible, but requires considerable effort) and *everlasting channels*, on which the information is intentionally published and recorded permanently (e.g. web pages that serve as a public bulletin board). Formally, we model everlasting privacy in the applied-pi calculus [1], a framework well-adapted to security protocols and already used to define privacy [11] and verifiability [16]. Our definitions apply not only to voting protocols but also to situations where forward secrecy is desirable, such as for instance untraceability in RFID protocols.

Modeling everlasting privacy also requires to precisely model what an attacker can break in the future. Our second contribution is a characterization, for several primitives, of what can be broken. The first natural primitive is encryption, for which we provide an equational theory that models the fact that private keys can be retrieved from public keys, or even from ciphertexts. Some other primitives have been primarily designed to achieve everlasting privacy. This is the case of *perfectly hiding* and *computationally binding* primitives, such as Pedersen commitments [19]. Intuitively, perfectly hiding means that the hidden secret cannot be retrieved even by a computationally unbounded adversary, while computationally binding means that, binding is ensured only for a (polynomially) limited attacker. We provide an equational theory that models such perfectly hiding and computationally binding primitives in general.

As an application, we study everlasting privacy for several variants of Helios [2], an e-voting protocol used for electing the president of the University of Louvain and board members of the IACR¹. We study in particular its latest variants with Pedersen

¹ International Association for Cryptologic Research.

commitments [12], designed to achieve everlasting privacy, still providing full verifiability. We also model and prove everlasting privacy of a (simplified) version of Moran and Naor’s protocol [18]. Interestingly, we were able to adapt algorithms in existing tools to automate the verification of everlasting privacy and we use adapted versions of the AKisS [6] and ProVerif [4] tools to analyze everlasting privacy for half a dozen of protocols.

Outline. In the following section we recall the applied pi calculus and introduce notations and basic definitions. In Section 3 we define new equivalence relations, namely forward and everlasting indistinguishability. Then, in Section 4 we instantiate these equivalences to the case of voting protocols, define everlasting privacy and illustrate this property on several examples. In Section 5 we present a modeling of perfectly hiding and computationally binding (and vice-versa) primitives in the applied pi calculus. In particular we model Pedersen commitments, which are for studying two protocols that provide everlasting privacy. In Section 6 we discuss tool support for automatically proving everlasting indistinguishability before concluding.

2 The Applied Pi Calculus

The applied pi calculus [1] is a language for modeling distributed systems and their interactions. It extends the pi calculus with an equational theory, which is particularly useful for modeling cryptographic protocols. The following subsections describe the syntax and semantics of the calculus.

2.1 Syntax

Terms. The calculus assumes an infinite set of names $\mathcal{N} = \{a, b, c, \dots\}$, an infinite set of variables $\mathcal{V} = \{x, y, z, \dots\}$ and a finite signature Σ , that is, a finite set of function symbols each with an associated arity. We use meta-variables u, v, w to range over both names and variables. Terms $M, N, T, \dots$ are built by applying function symbols to names, variables and other terms. Tuples $M_1, \dots, M_l$ are occasionally abbreviated $\vec{M}$. We write $\{M_1/u_1, \dots, M_l/u_l\}$ for substitutions that replace $u_1, \dots, u_l$ with $M_1, \dots, M_l$. The applied pi calculus relies on a simple type system. Terms can be of sort Channel for channel names or Base for the payload sent out on these channels. Function symbols can only be applied to, and return, terms of sort Base. A term is ground when it does not contain variables.

The signature Σ is equipped with an equational theory E , that is a finite set of equations of the form $M = N$. We define $=_E$ as the smallest equivalence relation on terms, that contains E and is closed under application of function symbols, substitution of terms for variables and bijective renaming of names.

Example 1. A standard signature for pairing and encryption is:

$$\Sigma_{\text{enc}} = \{0, 1, \langle _, _ \rangle, \text{fst}(_), \text{snd}(_), \text{pk}(_), \text{aenc}(_, _, _), \text{adec}(_, _), \text{senc}(_, _, _), \text{sdec}(_, _)\}$$

The term $\langle m_1, m_2 \rangle$ represents the concatenation of m_1 and m_2 , with associated projectors $\text{fst}(_)$ and $\text{snd}(_)$. The term $\text{aenc}(k, r, m)$ represents the asymmetric encryption of

$P, Q, R ::=$	processes
0	null process
$P \mid Q$	parallel composition
$!P$	replication
$\nu n.P$	name restriction
$u(x).P$	message input
$\overline{u}\langle M \rangle.P$	message output
$\text{if } M = N \text{ then } P \text{ else } Q$	conditional
$A, B, C ::=$	extended processes
P	plain process
$A \mid B$	parallel composition
$\nu n.A$	name restriction
$\nu x.A$	variable restriction
$\{M/x\}$	active substitution

where u is either a name or variable of channel sort.

Fig. 1. Applied pi calculus grammar

message m with public key k and randomness r while the associated decryption operator is adec . Similarly, $\text{senc}(k, r, m)$ represents the symmetric encryption of message m with key k and randomness r . The associated decryption operator is sdec . The properties of these primitives are modeled by the following standard equational theory E_{enc} :

$$E_{\text{enc}} = \left\{ \begin{array}{l} \text{fst}(\langle x, y \rangle) = x \\ \text{snd}(\langle x, y \rangle) = y \\ \text{adec}(x, \text{aenc}(\text{pk}(x), y, z)) = z \\ \text{sdec}(x, \text{senc}(x, y, z)) = z \end{array} \right\}$$

Processes. The grammar for processes is shown in Figure 1. Plain processes are standard. Extended processes introduce *active substitutions* which generalize the classical let construct: the process $\nu x.(\{M/x\} \mid P)$ corresponds exactly to the process let $x = M$ in P . As usual names and variables have scopes which are delimited by restrictions and by inputs. All substitutions are assumed to be cycle-free.

The sets of free and bound names, respectively variables, in process A are denoted by $\text{fn}(A)$, $\text{bn}(A)$, $\text{fv}(A)$, $\text{bv}(A)$. We also write $\text{fn}(M)$, $\text{fv}(M)$ for the names, respectively variables, in term M . An extended process A is *closed* if it has no free variables. A *context* $C[-]$ is an extended process with a hole. We obtain $C[A]$ as the result of filling $C[-]$'s hole with A . An *evaluation context* is a context whose hole is not under a replication, a conditional, an input, or an output.

Example 2. Throughout the paper we illustrate our definitions with a simplified version of the Helios voting system [2]. Two techniques can be used for tallying in Helios: either a homomorphic tally based on El Gamal encryption, or a tally based on mixnets. We present here the version with mixnets.

1. The voter V computes her ballot by encrypting her vote with the public key $\text{pk}(skE)$ of the election. The corresponding secret key is shared among several election authorities. Then she casts her ballot together with her identity on an authenticated channel. Upon receiving the ballot, the administrator simply publishes it on a public web page (after having checked that V is entitled to vote).
2. Once the voting phase is over, the votes are shuffled and reencrypted through mixnets. The permuted and rerandomized votes are again published on the public web page (together with a zero knowledge proof of correct reencryption and mixing).
3. Finally, the authorities decrypt the rerandomized votes and the administrator publishes the decrypted votes (with a zero knowledge proof of correct decryption).

The process representing the voter is parametrized by her vote v , and her identity id .

$$V(auth, id, v) \stackrel{\text{def}}{=} \nu r. \overline{auth} \langle \langle id, \text{aenc}(\text{pk}(skE), r, v) \rangle \rangle$$

The administrator BB receives votes on private authenticated channels and publishes the votes. It is parametrized by the authenticated channels of the voters. Then the ballots are forwarded to the tally T over the private channel c . The tally consists in decrypting the vote. The shuffle through mixnets is modeled simply, by non deterministic parallel composition after all ballots have been received. For simplicity, we consider here an election system for three voters.

$$\begin{aligned} BB(a_1, a_2, a_3) &\stackrel{\text{def}}{=} \nu c. a_1(x). \overline{bb} \langle x \rangle. \overline{c} \langle x \rangle \mid a_2(y). \overline{bb} \langle y \rangle. \overline{c} \langle y \rangle \mid a_3(z). \overline{bb} \langle z \rangle. \overline{c} \langle z \rangle \mid T \\ T &\stackrel{\text{def}}{=} c(x'). c(y'). c(z'). \\ &\quad (\overline{bb} \langle \text{adec}(skE, \text{snd}(x')) \rangle \mid \overline{bb} \langle \text{adec}(skE, \text{snd}(y')) \rangle \mid \overline{bb} \langle \text{adec}(skE, \text{snd}(z')) \rangle) \end{aligned}$$

The process H then represents the whole Helios system with two honest voters and one dishonest voter (which does therefore not need to be explicitly specified and whose authenticated channel is public).

$$H \stackrel{\text{def}}{=} \nu skE. \nu auth_1. \nu auth_2.$$

$$\overline{bb} \langle \text{pk}(skE) \rangle. (V(auth_1, id_1, a) \mid V(auth_2, id_2, b) \mid BB(auth_1, auth_2, auth_3))$$

The first honest voter casts the vote a while the second honest voter casts the vote b .

2.2 Semantics

The operational semantics of the applied pi calculus is defined by the means of two relations: structural equivalence and internal reductions. *Structural equivalence* ($\equiv$) is the smallest equivalence relation closed under α -conversion of both bound names and variables and application of evaluation contexts such that:

$$\begin{array}{ll} A \mid 0 \equiv A & \nu n. 0 \equiv 0 \\ A \mid (B \mid C) \equiv (A \mid B) \mid C & \nu u. \nu v. A \equiv \nu v. \nu u. A \\ A \mid B \equiv B \mid A & A \mid \nu u. B \equiv \nu u. (A \mid B) \\ !P \equiv P \mid !P & \text{if } u \notin \text{fn}(A) \cup \text{fv}(A) \\ \nu x. \{M/x\} \equiv 0 & \{M/x\} \equiv \{N/x\} \\ \{M/x\} \mid A \equiv \{M/x\} \mid A\{M/x\} & \text{if } M =_E N \end{array}$$

$$\begin{array}{c}
a(x).P \xrightarrow{a(M)} P\{M/x\} \qquad \frac{A \xrightarrow{\alpha} A' \quad u \text{ does not occur in } \alpha}{\nu u.A \xrightarrow{\alpha} \nu u.A'} \\
\\
\overline{a}\langle u \rangle.P \xrightarrow{\overline{a}\langle u \rangle} P \qquad \frac{A \xrightarrow{\alpha} A' \quad \text{bv}(\alpha) \cap \text{fv}(B) = \text{bn}(\alpha) \cap \text{fn}(B) = \emptyset}{A \mid B \xrightarrow{\alpha} A' \mid B} \\
\\
\frac{A \xrightarrow{\overline{a}\langle u \rangle} A' \quad u \neq a}{\nu u.A \xrightarrow{\nu u.\overline{a}\langle u \rangle} A'} \qquad \frac{A \equiv B \quad B \xrightarrow{\alpha} B' \quad A' \equiv B'}{A \xrightarrow{\alpha} A'}
\end{array}$$

Fig. 2. Labelled reductions

Internal reduction ($\rightarrow$) is the smallest relation closed under structural equivalence, application of evaluation contexts satisfying the following rules:

$$\begin{array}{ll}
\text{COMM} & \overline{c}\langle x \rangle.P \mid c(x).Q \rightarrow P \mid Q \\
\text{THEN} & \text{if } N = N \text{ then } P \text{ else } Q \rightarrow P \\
\text{ELSE} & \text{if } L = M \text{ then } P \text{ else } Q \rightarrow Q \\
& \text{for ground terms } L, M \text{ where } L \neq_E M
\end{array}$$

Labelled reduction ($\xrightarrow{\alpha}$) extends the internal reduction and enables the environment to interact with the processes as defined in Figure 2. The label α is either an input, or the output of a channel name or a variable of base type.

We write $\Rightarrow$ for an arbitrary (possibly zero) number of internal reductions and $\xRightarrow{\alpha}$ for $\Rightarrow \xrightarrow{\alpha} \Rightarrow$. Whenever the equational theory is not clear from the context we annotate the above relations by the equational theory and write e.g. $\rightarrow_E$.

A *trace* of a process is the sequence of actions (i.e. labels) together with the corresponding sent messages. Formally, the set of traces of a process A is defined as follows. Note that it depends on the underlying equational theory E .

$$\text{trace}_E(A) = \{(\alpha_1 \cdot \alpha_2 \cdot \dots \cdot \alpha_n, \varphi(B)) \mid A \xRightarrow{\alpha_1}_E A_1 \xRightarrow{\alpha_2}_E \dots A_{n-1} \xRightarrow{\alpha_n}_E B\}$$

Example 3. Consider the process H representing the Helios protocol as defined in Example 2. A possible execution for H is:

$$\begin{array}{c}
H \xrightarrow{\nu xk. \overline{bb}\langle xk \rangle} H_1 \\
\begin{array}{c}
\xrightarrow{\nu x. \overline{bb}\langle x \rangle} \xrightarrow{\nu y. \overline{bb}\langle y \rangle} \xrightarrow{\text{auth}_3(\langle id_3, x \rangle)} \xrightarrow{\nu z. \overline{bb}\langle z \rangle} \xrightarrow{\nu x'. \overline{bb}\langle x' \rangle} \xrightarrow{\nu y'. \overline{bb}\langle y' \rangle} \xrightarrow{\nu z'. \overline{bb}\langle z' \rangle} \\
H_2
\end{array}
\end{array}$$

where H_1 and H_2 are defined below (we omit the other intermediate processes). Note that H_2 is simply an active substitution.

$$\begin{array}{l}
H_1 = \nu skE. \nu auth_1. \nu auth_2. \nu r_1. \\
(\{\text{pk}(skE)/xk\} \mid V(auth_1, id_1, a) \mid V(auth_2, id_2, b) \mid BB(auth_1, auth_2, auth_3))
\end{array}$$

$$\begin{array}{l}
H_2 = \nu skE. \nu auth_1. \nu auth_2. \nu r_1. \nu r_2. \{\text{pk}(skE)/xk\} \mid \{a/x', b/y', a/z'\} \mid \\
\{\langle id_1, \text{aenc}(\text{pk}(skE), r_1, a) \rangle/x, \langle id_2, \text{aenc}(\text{pk}(skE), r_2, b) \rangle/y, \langle id_3, \text{aenc}(\text{pk}(skE), r_1, a) \rangle/z\}
\end{array}$$

This execution trace corresponds to the case where the two honest voters cast their vote as expected, while the dishonest voter replays the first voter's ballot. As we shall see in Example 5, this corresponds to the attack on privacy discovered in [9].

2.3 Equivalence Relations for Processes

Privacy is often stated in terms of equivalence [11]. We recall here the definitions of static and trace equivalence.

Sequences of messages are often stored as *frames*. Formally, a frame is an extended process built from 0 and active substitutions $\{M/x\}$, and closed by parallel composition and restriction. The *domain* of a frame $\phi = \nu \tilde{n}. \{M_1/x_1, \dots, M_n/x_n\}$ such that $x_i \notin \tilde{n}$ is $\text{dom}(\phi) = \{x_1, \dots, x_n\}$. Every extended process A can be mapped to a frame $\varphi(A)$ by replacing every plain process in A with 0. The frame $\varphi(A)$ represents the static knowledge output by a process to its environment.

Two frames are indistinguishable to an attacker if it is impossible to build a test that allows to differentiate between the two.

Definition 1 (Static equivalence). *Given an equational theory E two frames ϕ and ψ are statically equivalent, denoted $\phi \sim_E \psi$, if $\text{dom}(\phi) = \text{dom}(\psi)$ and there exist $\tilde{n}, \sigma, \tau$ such that $\phi \equiv \nu \tilde{n}. \sigma$, $\psi \equiv \nu \tilde{n}. \tau$ and for all terms M, N such that $\tilde{n} \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$, we have $M\sigma \equiv_E N\sigma$ if and only if $M\tau \equiv_E N\tau$. By abuse of notation, we may write $M\phi$ instead of $M\sigma$ when σ is clear from the context.*

Example 4. Let E_{enc} be the equational theory defined at Example 1. Let H_2 be the process/frame defined in Example 3. Let $\phi = \varphi(H_2)$ ($= H_2$ actually). Consider the following frame ψ .

$$\psi = \nu skE. \nu r_1. \nu r_2. \{\text{pk}(skE)/x_k\} \mid \{a/x', b/y', b/z'\} \mid \{\langle id_1, \text{aenc}(\text{pk}(skE), r_1, b) \rangle / x, \langle id_2, \text{aenc}(\text{pk}(skE), r_2, a) \rangle / y, \langle id_3, \text{aenc}(\text{pk}(skE), r_1, b) \rangle / z, \}$$

The two frames ϕ and ψ are not statically equivalent for the equational theory E_{enc} . Indeed, consider for example $M = z'$ and $N = a$, we have $M\phi = a = N\phi$ but $M\psi = b \neq N\psi$. Therefore, $\phi \not\sim_{E_{\text{enc}}} \psi$.

The active counterpart of static equivalence is trace equivalence. Intuitively, two processes A and B are indistinguishable to an attacker if any execution of A can be matched to an execution of B that is equal for their observable actions and such that the corresponding sequences of sent messages are statically equivalent.

Definition 2 (Trace equivalence). *Given an equational theory E two processes A and B are trace equivalent, denoted $A \stackrel{\text{tr}}{\sim}_E B$, if for any trace $(tr_A, \phi_A) \in \text{trace}_E(A)$ there is a corresponding trace $(tr_B, \phi_B) \in \text{trace}_E(B)$ such that $tr_A = tr_B$ and $\phi_A \sim_E \phi_B$ (and reciprocally).*

Example 5. We consider the Helios system H' with two honest voters and one dishonest voter where one honest voter casts the vote b while the other one casts the vote a .

$$H' \stackrel{\text{def}}{=} \nu skE. \nu auth_1. \nu auth_2.$$

$$\overline{bb}\langle pk(skE) \rangle. (V(auth_1, id_1, b) \mid V(auth_2, id_2, a) \mid BB(auth_1, auth_2, auth_3))$$

Let (tr, ϕ) be the trace corresponding to the execution of H described in Example 3 where $\phi = \varphi(H_2) = H_2$ (as defined in Example 3) and $tr = \nu xk. \overline{bb}\langle xk \rangle \cdot \nu x. \overline{bb}\langle x \rangle \cdot \nu y. \overline{bb}\langle y \rangle \cdot auth_3(\langle id_3, x \rangle) \cdot \nu z. \overline{bb}\langle z \rangle \cdot \nu x'. \overline{bb}\langle x' \rangle \cdot \nu y'. \overline{bb}\langle y' \rangle \cdot \nu z'. \overline{bb}\langle z' \rangle$. Then $(tr, \phi) \in \text{trace}_{\text{Enc}}(H)$ and for any $(tr, \phi') \in \text{trace}_{\text{Enc}}(H')$, it is easy to check that we have $\phi \not\sim_{\text{Enc}} \phi'$. (In fact, $\phi' = \psi$ from Example 4.) Therefore, $H \not\sim_{\text{Enc}}^{\text{tr}} H'$

Intuitively, if the dishonest voter's strategy is to replay the first voter's vote, then he would cast a vote of the form $\langle id_3, \text{aenc}(pk(skE), r_1, a) \rangle$ in the system H while he would cast a vote of the form $\langle id_3, \text{aenc}(pk(skE), r_1, b) \rangle$ in the system H' . Once the result is published, an attacker would be then able to distinguish H from H' since the tally in H is $\{a, b, a\}$ while it is $\{b, a, b\}$ in H' . This corresponds exactly to the replay attack against Helios, explained in [9].

3 Forward and Everlasting Indistinguishability

In this section we introduce and illustrate our definitions of forward and everlasting indistinguishability. In the next section we will show how the here presented definitions can be used to define practical everlasting privacy in electronic voting.

3.1 Definitions of Forward and Everlasting Indistinguishability

From now on we suppose that Σ is a signature and that E_0 and E_1 are equational theories over Σ . We want to model that an attacker may interact with a protocol today and store some data which may be exploited in the future when his computational power has increased. We model the fact that the attacker's computational power may change by using two different equational theories: E_0 models the attacker's capabilities while interacting with the protocol at the time of the election, while E_1 models his capabilities when exploiting the published data in the future when the strength of cryptography has been eroded.

We also argue that in many situations it is reasonable to suppose that the attacker does not store all of the data that was sent over the network. We will therefore consider some channels to be *everlasting*: data sent over such channels will remain in the attacker's knowledge for future analysis while other data will be “forgotten” and can only be used during the interaction with the protocol. Typically, everlasting channels are media such as web-pages published on the Internet (that can easily be accessed by anyone, for a rather long period of time) while public but non-everlasting channels can be communications over the Internet, which can be recorded only by the active and costly involvement of an attacker.

In order to reason about data that has been sent on certain channels we introduce the following notation. Let P be a process, $\mathcal{C}$ a set of channels (corresponding to the

everlasting channels), and $tr = (\alpha_1 \cdot \alpha_2 \cdot \dots \cdot \alpha_n, \varphi) \in \text{trace}_E(P)$ a trace of P . We define the set of variables in the domain of φ corresponding to terms sent on channels in $\mathcal{C}$ as $\mathcal{V}_{\mathcal{C}}(\alpha_1 \cdot \alpha_2 \cdot \dots \cdot \alpha_n) = \{x \mid c \in \mathcal{C}, 1 \leq i \leq n, \alpha_i = \nu x. \bar{c}\langle x \rangle\}$ and denote by $\phi_{\mathcal{V}}(P_n)$ the substitution $\phi(P_n)$ whose domain is restricted to the set of variables $\mathcal{V}$.

Two processes A and B are said to be forward indistinguishable if, informally, an attacker cannot observe the difference between A and B being given the computational power modeled by E_1 (where it can break keys for example), but for executions that happened in the past, that is over E_0 (the standard theory) and observing only the information that was passed through everlasting channels.

Definition 3 (Forward indistinguishability). *Let A and B be two closed extended processes and $\mathcal{C}$ a set of channels. We define $A \sqsubseteq_{E_0, E_1}^{\text{fwd}, \mathcal{C}} B$, if for every trace $(\alpha_1 \cdot \alpha_2 \cdot \dots \cdot \alpha_n, \varphi_A) \in \text{trace}_{E_0}(A)$ there exists φ_B s.t. $(\alpha_1 \cdot \alpha_2 \cdot \dots \cdot \alpha_n, \varphi_B) \in \text{trace}_{E_0}(B)$*

$$\text{and } \phi_{A\mathcal{V}} \sim_{E_1} \phi_{B\mathcal{V}}.$$

where $\mathcal{V} = \mathcal{V}_{\mathcal{C}}(\alpha_1 \cdot \alpha_2 \cdot \dots \cdot \alpha_n)$. A and B are forward indistinguishable w.r.t. $\mathcal{C}$, E_0 and E_1 , denoted $A \approx_{E_0, E_1}^{\text{fwd}, \mathcal{C}} B$, if $A \sqsubseteq_{E_0, E_1}^{\text{fwd}, \mathcal{C}} B$ and $B \sqsubseteq_{E_0, E_1}^{\text{fwd}, \mathcal{C}} A$.

Note that in the above definition we only check equivalence of messages that were sent on channels in the set $\mathcal{C}$. We may also require that A and B are indistinguishable in the standard way (over E_0). Standard indistinguishability and forward indistinguishability yield *everlasting indistinguishability*.

Definition 4 (Everlasting indistinguishability). *Let A and B be two closed extended processes, $\mathcal{C}$ a set of channels. A and B are everlasting indistinguishable w.r.t. $\mathcal{C}$, E_0 and E_1 , denoted $A \approx_{E_0, E_1}^{\text{ev}, \mathcal{C}} B$ if*

1. $A \approx_{E_0}^{\text{tr}} B$, i.e. A and B are trace equivalent w.r.t. E_0 ; and
2. $A \approx_{E_0, E_1}^{\text{fwd}, \mathcal{C}} B$, i.e. A and B are forward indistinguishable w.r.t. $\mathcal{C}$, E_0 and E_1 .

3.2 Examples

We illustrate the above definitions on a simple RFID protocol. In the context of RFID systems, forward secrecy is often a desired property: even if an RFID tag has been tampered with, and its secrets have been retrieved by a malicious entity, its past transactions should remain private. This can be seen as a form of everlasting security requirement. Indeed, RFID tags being devices vulnerable to tampering, one would like to make sure that when an intruder gains access to an honest device, he is not able to trace back the movements of the tag. Such tampering can be modelled by the following equational theory E_{break} , that gives direct access to keys.

$$E_{\text{break}} = \left\{ \begin{array}{l} \text{break}_{\text{aenc}}(\text{aenc}(\text{pk}(x), y, z)) = x \\ \text{break}_{\text{senc}}(\text{senc}(x, y, z)) = x \end{array} \right\}$$

We also use this equational theory later to model that in 20 or 30 years an adversary will be able to break nowadays encryption keys.

Consider the following toy RFID protocol

$$\begin{aligned} \text{session} &= \nu r. \bar{c}(\text{enc}(k, r, id)) \\ \text{tag} &= \nu k. \nu id. !\text{session} \\ \text{system} &= !\text{tag} \end{aligned}$$

where a tag identifies itself to a reader by sending its tag identifier id encrypted with a long-term symmetric key shared between the tag and the reader.

We can model unlinkability as being the property that an attacker cannot distinguish the situation where the same tag is used in several sessions from the situation where different tags are used. Formally unlinkability is modelled as the trace equivalence:

$$\text{system} \approx_{\text{E}_{\text{enc}}}^{\text{tr}} \text{system}'$$

where

$$\text{system}' = !\nu k. \nu id. \text{session}.$$

Intuitively, this protocol satisfies unlinkability only as long as the keys are not leaked. Indeed, since each identification uses a different random in the encrypted message sent to the reader, each of the sent messages is different and looks like a random message to the intruder. However, system and system' are not forward indistinguishable when considering a theory E_1 which allows to break keys, i.e.,

$$\text{system} \not\approx_{\text{E}_{\text{enc}}, \text{E}_{\text{enc}} \cup \text{E}_{\text{break}}}^{\text{fwd}\{c\}} \text{system}'$$

where E_{enc} is the equational theory introduced in Example 1. Indeed, once the key k of a tag is obtained by the intruder, he can retrieve the identity behind each blob he has seen on channel c , and thus distinguish the set of messages obtained by an execution of system where the same tag executes at least two sessions, from the set of messages obtained by the corresponding execution of system' where each tag has executed at most one session. Thus this protocol does not satisfy the stronger requirement of everlasting indistinguishability either:

$$\text{system} \not\approx_{\text{E}_{\text{enc}}, \text{E}_{\text{enc}} \cup \text{E}_{\text{break}}}^{\text{ev}\{c\}} \text{system}'$$

4 Application to Practical Everlasting Privacy

We model a practical version of everlasting privacy in voting protocols based on everlasting indistinguishability.

4.1 Definition of Practical Everlasting Privacy

We first recall the definition of vote privacy introduced in [15].

Definition 5 (Vote privacy). Let $\text{VP}(v_1, v_2)$ be an extended process with two free variables v_1, v_2 . $\text{VP}(v_1, v_2)$ respects vote privacy for an equational theory E if

$$\text{VP}(a, b) \approx_{\text{E}}^{\text{tr}} \text{VP}(b, a)$$

Intuitively, the free variables refer to the votes of two honest voters id_1 and id_2 . Then this equivalence ensures that an attacker cannot distinguish the situations where voter id_1 voted for candidate a and voter id_2 voted for candidate b , from the situation where the voters swapped their votes, i.e., voter id_1 voted for candidate b and voter id_2 voted for candidate a .

Example 6. Let $\text{Helios}(v_1, v_2)$ be the process

$$\nu skE. \nu auth_1. \nu auth_2. \\ \overline{bb}\langle pk(skE) \rangle. (V(auth_1, id_1, v_1) \mid V(auth_2, id_2, v_2) \mid BB(auth_1, auth_2, auth_3))$$

where processes V and BB are defined in Example 2.

In Example 5, when we illustrated trace equivalence we showed that Helios does not satisfy vote privacy due to a vote replay attack discovered in [9].

A simple fix of the attack consists in weeding duplicates. The corresponding tally is

$$T' \stackrel{\text{def}}{=} c(x').c(y').c(z'). \\ \text{if } \text{snd}(x') \neq \text{snd}(y') \wedge \text{snd}(x') \neq \text{snd}(z') \wedge \text{snd}(y') \neq \text{snd}(z') \text{ then} \\ \overline{bb}\langle \text{adec}(skE, \text{snd}(x')) \rangle \mid \overline{bb}\langle \text{adec}(skE, \text{snd}(y')) \rangle \mid \overline{bb}\langle \text{adec}(skE, \text{snd}(z')) \rangle$$

In other words, the tally is performed only if there are no duplicates amongst the cast votes. We define the voting protocol $\text{Helios}^{\text{noreplay}}$ as Helios but with the revised version T' of the tally. Using the tools ProVerif and AKISS we have shown that this protocol satisfies vote privacy.

$$\text{Helios}^{\text{noreplay}}(a, b) \stackrel{\text{tr}}{\approx}_{E_{\text{enc}}} \text{Helios}^{\text{noreplay}}(b, a)$$

The above definition of vote privacy does however not take into account that most cryptographic schemes rely on computational assumptions and may be broken in the future. In order to protect the secrecy of votes against such attacks in the future we propose a stronger definition based on forward indistinguishability.

Definition 6 (Everlasting vote privacy). Let $\text{VP}(v_1, v_2)$ be an extended process with two free variables v_1, v_2 . $\text{VP}(v_1, v_2)$ satisfies everlasting privacy w.r.t. a set of channels $\mathcal{C}$ and equational theories E_0 and E_1 , if

$$\text{VP}(a, b) \stackrel{\text{ev } \mathcal{C}}{\approx}_{E_0, E_1} \text{VP}(b, a)$$

We note that everlasting vote privacy is strictly stronger than vote privacy as it requires trace equivalence of $\text{VP}(a, b)$ and $\text{VP}(b, a)$ (which is exactly vote privacy) and additionally forward indistinguishability of these processes. Our definition is parametric with respect to the equational theories and the channels we suppose to be everlasting. The equational theory E_1 allows us to exactly specify what a future attacker may be able to break. The set of everlasting channels $\mathcal{C}$ allows us to specify what data a future attacker has access to. When $\mathcal{C}$ corresponds to all channels we typically get a requirement which is too strong for practical purposes. We argue that it is reasonable to suppose that in, say 50 years, an attacker does not have access to the transmissions between individual voters and the system while a bulletin board published on the Internet could easily have been stored.

4.2 Examples

Helios with Identities. As discussed In Example 6, $\text{Helios}^{\text{noreplay}}$ does satisfy vote privacy. However, this protocol does not satisfy everlasting vote privacy with $E_0 = E_{\text{enc}}$, $E_1 = E_{\text{enc}} \cup E_{\text{break}}$ and $C = \{bb\}$. Intuitively, this is due to the fact that a future attacker can break encryption and link the recovered vote to the identity submitted together with the cast ballot. Formally, we can show that

$$\text{Helios}^{\text{noreplay}}(a, b) \stackrel{\text{fwd}}{\not\sim} \text{Helios}^{\text{noreplay}}(b, a)$$

Consider the trace $(\nu xk. \overline{bb}\langle xk \rangle \cdot \nu x. \overline{bb}\langle x \rangle \cdot \nu y. \overline{bb}\langle y \rangle, \varphi_A) \in \text{trace}_{E_{\text{enc}}}(\text{Helios}^{\text{noreplay}}(a, b))$ where

$$\varphi_A = \nu skE, r_1, r_2. \{ \text{pk}(skE)/xk, \\ \langle id_1, \text{aenc}(\text{pk}(skE), r_1, a) \rangle /x, \\ \langle id_2, \text{aenc}(\text{pk}(skE), r_2, b) \rangle /y \}$$

Traces $(\nu xk. \overline{bb}\langle xk \rangle \cdot \nu x. \overline{bb}\langle x \rangle \cdot \nu y. \overline{bb}\langle y \rangle, \varphi_B) \in \text{trace}_{E_{\text{enc}}}(\text{Helios}^{\text{noreplay}}(b, a))$ are either such that

$$\varphi_B \equiv \nu skE, r_1, r_2. \{ \text{pk}(skE)/xk, \\ \langle id_1, \text{aenc}(\text{pk}(skE), r_1, b) \rangle /x, \\ \langle id_2, \text{aenc}(\text{pk}(skE), r_2, a) \rangle /y \}$$

or

$$\varphi_B \equiv \nu skE, r_1, r_2. \{ \text{pk}(skE)/xk, \\ \langle id_2, \text{aenc}(\text{pk}(skE), r_1, a) \rangle /x, \\ \langle id_1, \text{aenc}(\text{pk}(skE), r_2, b) \rangle /y \}$$

In both cases we have that $\varphi_A \not\sim_{E_{\text{enc}} \cup E_{\text{break}}} \varphi_B$. In the first case this is witnessed by the test $M = a$ and $N = \text{break}_{\text{aenc}}(\text{snd}(x))$ as

$$M\varphi_A = a =_{E_{\text{enc}} \cup E_{\text{break}}} N\varphi_A \quad \text{but} \quad M\varphi_B = a \neq_{E_{\text{enc}} \cup E_{\text{break}}} b =_{E_{\text{enc}} \cup E_{\text{break}}} N\varphi_B$$

In the second case non equivalence is witnessed by the test $M = id_1$ and $N = \text{fst}(x)$.

Helios without Identities. As we just saw $\text{Helios}^{\text{noreplay}}$ does not satisfy everlasting privacy. This is due to the fact that encrypted votes are published together with the identity of the voter on the bulletin board. A simple variant (used e.g. in Louvain for student elections) consists in publishing the encrypted vote without the identity of the voter. We define $\text{Helios}^{\text{noide}}$ as $\text{Helios}^{\text{noreplay}}$ but redefining the process BB' as

$$BB'(a_1, a_2, a_3) \stackrel{\text{def}}{=} \nu c. a_1(x). \overline{bb}\langle \text{snd}(x) \rangle. \overline{c}\langle x \rangle \mid a_2(y). \overline{bb}\langle \text{snd}(y) \rangle. \overline{c}\langle y \rangle \mid \\ a_3(z). \overline{bb}\langle \text{snd}(z) \rangle. \overline{c}\langle z \rangle \mid T'$$

where T' is as defined in Example 6. As we shall see in Section 6, we prove everlasting privacy of $\text{Helios}^{\text{noide}}$ w.r.t E_{enc} , E_{break} and everlasting channel bb , using (adaptations of) ProVerif and AKISS.

5 Modeling Commitments

Commitment schemes allow a sender to commit to a value v while keeping this value hidden until an ‘opening’ phase, where the sender reveals v to the receiver of the commitment $\text{commit}(v)$. The receiver should then be able to verify that the revealed value is indeed the one used to compute $\text{commit}(v)$, and in that sense that the sender had indeed committed to the revealed value. The two main security properties of such schemes are *binding* (the sender can’t claim that $\text{commit}(v)$ is a commitment to some $v' \neq v$), and *hiding* (the receiver can’t deduce v from $\text{commit}(v)$). These two properties can hold ‘perfectly’ or merely ‘computationally’. It is known that there are no commitment schemes which are both perfectly hiding and perfectly binding, so one has to choose between perfectly hiding and computationally binding (PHCB) and perfectly binding and computationally hiding (PBCH). In this section, we characterize in our formal model what it means for a primitive to be PHCB and PBCH. We also give equational theories to model such primitives, which we then use for the verification of two voting protocols that rely on such primitives to ensure everlasting vote privacy.

5.1 Modeling Hiding and Binding Cryptographic Primitives

PBCH Primitives. Informally, an n -ary function f is *perfectly binding* if the inputs are totally determined by the output. In other words, f is perfectly binding if it admits no collisions. It is *computationally hiding* if it is hard to retrieve the inputs from the output.

To model a PBCH primitive f using the applied pi calculus, we introduce two equational theories E_0^f and E_1^f over the signature $\Sigma = \{f, \text{break}_f^1, \dots, \text{break}_f^n\}$, such that no equation of the form

$$f(M_1, \dots, M_n) =_E f(N_1, \dots, N_n)$$

is derivable, where $(M_1, \dots, M_n) \neq_E (N_1, \dots, N_n)$ and $E \in \{E_0^f, E_1^f\}$; and that the equation

$$\text{break}_f^i(f(v_1, \dots, v_n)) =_{E_1^f} v_i.$$

is derivable. As before, E_0^f models the capabilities of a computationally bounded attacker interacting with the protocol, while E_1^f models the capabilities of a computationally unbounded attacker in the future.

Example 7. A trivial example of a perfectly binding function is the identity function id . However, id is not hiding.

Example 8. An example of a PBCH primitive is the ElGamal public key derivation function. Given multiplicative cyclic group G of order q with generator g , to generate a private and public key pair Alice does the following:

1. she chooses at random her private key $sk \in \{1, \dots, q-1\}$,
2. she computes and publishes her public key $\text{pk}_{G,g,q}(sk) = g^{sk}$.

The secret key sk is totally determined by the public key $\text{pk}_{G,g,q}(sk) = g^{sk}$. It is however as hard to find sk from $\text{pk}_{G,g,q}(sk)$ as it is to compute discrete logarithms.

Thus, to reason about protocols relying on ElGamal encryption we consider the following equational theories over the signature $\{\text{aenc}_{G,g,q}, \text{adec}_{G,g,q}, \text{pk}_{G,g,q}, \text{break}_{\text{pk}_{G,g,q}}\}$ (we omit the subscripts G, g, q for readability):

$$\begin{aligned} \mathbb{E}_0^{\text{ElGamal}} &= \{\text{adec}(xk, \text{aenc}(\text{pk}(xk), xr, xm)) = xm\} \\ \mathbb{E}_1^{\text{ElGamal}} &= \left\{ \begin{array}{l} \text{adec}(xk, \text{aenc}(\text{pk}(xk), xr, xm)) = xm \\ \text{break}_{\text{pk}}(\text{pk}(xk)) = xk \end{array} \right\} \end{aligned}$$

The function $\text{pk}_{G,g,q}$ is PBCH. Note however that the encryption algorithm $\text{aenc}_{G,g,q}$ is not PBCH, since it is not perfectly binding. Indeed, given the parameters G, q , and g , to encrypt the message m with the public key g^{sk} , Alice would

1. pick a random $r \in \{0, \dots, q-1\}$ and compute $c_1 = g^r$;
2. compute the secret shared key $s = (g^{sk})^r$; and
3. compute $c_2 = m \cdot s$

The computed ciphertext is then $\text{aenc}(\text{pk}(sk), r, m) = (c_1, c_2) = (g^r, m \cdot (g^{sk})^r)$. Hence, for any public key $\text{pk}(sk') = g^{sk'}$, there exists a message $m' = m \cdot (g^{sk})^r / (g^{sk'})^r$ such that $\text{aenc}(\text{pk}(sk), r, m) = \text{aenc}(\text{pk}(sk'), r, m')$. Thus, ElGamal encryption is not perfectly binding.

PHCB Primitives. Informally, an n -ary function f is perfectly hiding if given the output, it is impossible to retrieve any of the inputs. So even enumerating all the possible inputs shouldn't allow one to retrieve the inputs from the output of the function. But this implies that f should admit collisions for each possible input. On the other hand, f is computationally binding if it is computationally not feasible to find such collisions.

Example 9. Any constant function $f(x_1, \dots, x_n) = c$ is obviously perfectly hiding but not computationally binding. The $\oplus$ function is also perfectly hiding since for all $z = x \oplus y$

- for all x' , we have that $y' = z \oplus x'$ is such that $x \oplus y = x' \oplus y'$; and
- for all y'' , we have that $x'' = z \oplus y''$ is such that $x \oplus y = x'' \oplus y''$.

However, it is not computationally binding since it is easy to compute x'' and y' .

Example 10. Pedersen commitments are PHCB. The Pedersen commitment over a cyclic group G of order q and two generators $h, g \in G$ such that $\log_g h$ is not known is the function $P_{h,g}^G(x, y) = h^x \cdot g^y \pmod{q}$. Pedersen commitments are perfectly hiding since for all $z = P_{h,g}^G(x, y)$,

- for all x' , we have that $y' = y + (x - x') \cdot \log_g h \pmod{q}$ is such that $P_{h,g}^G(x, y) = P_{h,g}^G(x', y')$;
- for all y'' , we have that $x'' = x + (y - y'') \cdot \log_h g \pmod{q}$ is such that $P_{h,g}^G(x, y) = P_{h,g}^G(x'', y'')$.

but they are computationally binding because finding these x'' and y' is as hard as computing discrete logarithms.

To reason about protocols relying on Pedersen commitments using the applied pi calculus, we introduce the function symbols $\text{forge}_{\text{Ped}}^1$, and $\text{forge}_{\text{Ped}}^2$ and the two following equational theories

$$\begin{aligned} E_0^{\text{Ped}} &= \emptyset \\ E_1^{\text{Ped}} &= \left\{ \begin{array}{l} \text{Ped}(\text{forge}_{\text{Ped}}^1(v, y'), y') = v \\ \text{Ped}(x', \text{forge}_{\text{Ped}}^2(v, x')) = v \\ \text{forge}_{\text{Ped}}^1(\text{Ped}(x, y), y) = x \\ \text{forge}_{\text{Ped}}^2(\text{Ped}(x, y), x) = y \\ \text{forge}_{\text{Ped}}^1(v, \text{forge}_{\text{Ped}}^2(v, x)) = x \\ \text{forge}_{\text{Ped}}^2(v, \text{forge}_{\text{Ped}}^1(v, y)) = y \end{array} \right\} \end{aligned}$$

For the first equation, suppose $v = \text{Ped}(x, y)$, and we have some y' ; then $\text{forge}_{\text{Ped}}^1$ allows us to compute a value $x' = \text{forge}_{\text{Ped}}^1(v, y')$ such that $v = \text{Ped}(x', y')$. The second equation is similar. The third and fourth equation allow us to recover one of the arguments, given that the other argument is known. In other words the third equation expresses that when forging $x' = \text{forge}_{\text{Ped}}^1(v, y)$ and $v = \text{Ped}(x, y)$ then we must have that $x' = x$, and similarly for the fourth equation. The fifth and sixth equations are also seen to be mathematically valid, given that $\text{forge}_{\text{Ped}}^1(v, y)$ and $\text{forge}_{\text{Ped}}^2(v, x)$ respectively model the terms $\log_g(v/h^y)$ and $\log_h(v/g^x)$.

5.2 Applications: Electronic Voting Protocols and Everlasting Privacy

Pedersen commitments have been used in several voting protocols for achieving everlasting privacy. In particular we study the protocol by Moran and Naor [18] and a recent version of Helios [12] based on Pedersen commitments.

Moran-Naor Protocol. Moran and Naor [18] designed a protocol to be used with voting machines in a polling station. The protocol aims to achieve both verifiability and everlasting privacy. From a high level point of view the protocol works as follows.

1. The voter enters his vote into the voting machine inside the voting booth. The machine then computes a Pedersen commitment to this vote and provides a zero knowledge proof to the voter that the computed value commits to the voter's choice. The commitment is then published on a bulletin board so that the voter can verify the presence of his ballot.
2. After all ballots have been cast, the votes are published (in random order) on the bulletin board together with a zero knowledge proof asserting that the published votes correspond to the votes of the published commitments.

As we are only interested in privacy and not verifiability we ignore the zero knowledge proofs in our modeling and simply represent the protocol by the process

$$\begin{aligned} \text{MoranNaor}(v_1, v_2) &\stackrel{\text{def}}{=} \nu \text{priv}_1. \nu \text{priv}_2. \\ &V(\text{priv}_1, v_1) \mid V(\text{priv}_2, v_2) \mid \nu c.(DRE(\text{priv}_1, \text{priv}_2, \text{priv}_3) \mid T) \end{aligned}$$

where

$$\begin{aligned}
 V(priv, v) &\stackrel{\text{def}}{=} \overline{priv}\langle v \rangle \\
 DRE(p_1, p_2, p_3) &\stackrel{\text{def}}{=} p_1(x_1). \nu r_1. \overline{bb}\langle \text{Ped}(x_1, r_1) \rangle. \overline{c}\langle x_1 \rangle \mid \\
 &\quad p_2(x_2). \nu r_2. \overline{bb}\langle \text{Ped}(x_2, r_2) \rangle. \overline{c}\langle x_2 \rangle \mid \\
 &\quad p_3(x_3). \nu r_3. \overline{bb}\langle \text{Ped}(x_3, r_3) \rangle. \overline{c}\langle x_3 \rangle \\
 T &= c(y_1). \overline{bb}\langle y_1 \rangle \mid c(y_2). \overline{bb}\langle y_2 \rangle \mid c(y_3). \overline{bb}\langle y_3 \rangle
 \end{aligned}$$

As the voter enters his vote in a private ballot booth, we have modelled this communication on a private channel. We have been able to show that MoranNaor verifies everlasting privacy with respect to the channel bb and the equational theories introduced for Pedersen commitments.

Helios with Pedersen Commitments. In [12], the authors propose a version of the Helios voting system that provides everlasting vote privacy *w.r.t.* the bulletin board. They rely for this on Pedersen commitments. In this section, we present this new version of the Helios system.

1. The voter V chooses her candidate v and commits to it by generating a random number r and computing the Pedersen commitment $\text{Ped}(r, v)$. She then separately encrypts the decommitment values r and v using the public key $\text{pk}(\text{sk}E)$ of the election; and casts her commitment together with the encrypted decommitment values and her identity on a private authenticated channel. Upon reception of the ballot, the Bulletin Board (BB) publishes on a public web page the commitment $\text{Ped}(r, v)$ (after having checked that V is entitled to vote).
2. Once the voting phase is over, the ballots (*i.e.* the commitments together with the encrypted decommitment values) are shuffled and rerandomized through mixnets. The random permutation of the rerandomized ballots is published on the public webpage (together with a zero knowledge proof of correct reencryption and mixing).
3. Finally, the authorities decrypt the rerandomized and shuffled decommitment values and the BB publishes them.

The voter can be modelled by the following process:

$$\begin{aligned}
 V(id, auth, v) &\stackrel{\text{def}}{=} \nu s. \nu rv. \nu rs. \\
 &\quad \overline{auth}\langle \langle id, \langle \text{Ped}(s, v), \langle \text{aenc}(\text{pk}(\text{sk}E), rv, v), \text{aenc}(\text{pk}(\text{sk}E), rs, s) \rangle \rangle \rangle \rangle
 \end{aligned}$$

She sends to the BB on the private authenticated channel $authCh$, her commitment $\text{Ped}(s, v)$ to vote v , together with her identity and the encrypted decommitment values $\text{aenc}(\text{pk}(\text{sk}E), rv, v), \text{aenc}(\text{pk}(\text{sk}E), rs, s)$.

The ballot box publishes her commitment for verifiability purposes. After having received all votes, the BB publishes the votes in a random order through the process T .

$$\begin{aligned}
 BB(a_1, a_2, a_3) &\stackrel{\text{def}}{=} a_1(x). \overline{bb}\langle \langle \text{fst}(x), \text{fst}(\text{snd}(x)) \rangle \rangle. \overline{c}\langle x \rangle \mid \\
 &\quad a_2(y). \overline{bb}\langle \langle \text{fst}(x), \text{fst}(\text{snd}(x)) \rangle \rangle. \overline{c}\langle x \rangle \mid \\
 &\quad a_3(z). \overline{c}\langle z \rangle \mid T
 \end{aligned}$$

$$\begin{aligned}
T \stackrel{\text{def}}{=} & c(x). c(y). c(z). \text{if fst}(\text{snd}(\text{snd}(x))) \neq \text{fst}(\text{snd}(\text{snd}(z))) \\
& \wedge \text{fst}(\text{snd}(\text{snd}(y))) \neq \text{fst}(\text{snd}(\text{snd}(z))) \\
& \wedge \text{fst}(\text{snd}(\text{snd}(x))) \neq \text{fst}(\text{snd}(\text{snd}(y))) \text{ then} \\
& \overline{bb}\langle \text{adec}(skE, \text{fst}(\text{snd}(\text{snd}(x)))) \rangle \mid \\
& \overline{bb}\langle \text{adec}(skE, \text{fst}(\text{snd}(\text{snd}(y)))) \rangle \mid \\
& \overline{bb}\langle \text{adec}(skE, \text{fst}(\text{snd}(\text{snd}(z)))) \rangle
\end{aligned}$$

Finally we can define the voting protocol $\text{Helios}^{\text{Ped}}$ as

$$\begin{aligned}
\text{Helios}^{\text{Ped}}(v1, v2) \stackrel{\text{def}}{=} & \nu skE. \nu auth_1. \nu auth_2. \\
& \overline{bb}\langle \text{pk}(skE) \rangle. (V(auth_1, id_1, v1) \mid V(auth_2, id_2, v2) \mid BB(auth_1, auth_2, auth_3))
\end{aligned}$$

which verifies everlasting privacy with respect to the channel bb and the previously introduced equational theories.

6 Tool Support for Everlasting Indistinguishability

In order to verify everlasting indistinguishability on the examples presented in the previous section we have adapted two tools for automated verification of equivalence properties, AKISS [6] and ProVerif [5]. The two tools have shown themselves to be complementary and the results obtained using the tools are summarized in Figure 3.

AKISS. AKISS is a recent tool that has been designed to automatically prove trace equivalence by translating processes into Horn clauses and using a dedicated resolution algorithm. More precisely it can both under- and over-approximate trace equivalence in the case of a bounded number of sessions, i.e. for processes without replication. The tool has currently two limitations: it does not support private channels, or else branches in conditionals. However, it is able to deal with a wide range of equational theories, including the theory for Pedersen commitments introduced in the previous section.

We have adapted the tool in order to check forward indistinguishability and adapted the syntax to declare everlasting channels and an everlasting equational theory. More precisely we implemented an algorithm to check an under-approximation of forward indistinguishability, yielding a proof of forward indistinguishability whenever the tool responds positively. While false attacks are possible, we did not encounter any in our case studies.

Absence of support for private channels and else branches required us to adapt some of the examples. In particular we rewrote the processes by directly *inlining* private communications, which in the examples maintained the same set of traces, hence preserving everlasting indistinguishability. The weeding operation in $\text{Helios}^{\text{noreplay}}$, $\text{Helios}^{\text{noid}}$ and $\text{Helios}^{\text{ped}}$ requires the use of an else branch. We encoded a different weeding procedure using cryptographic proofs of knowledge. While the vote replay attack on the simple Helios protocol is found in less than 10 seconds, the verification of other examples ranged from a few minutes to several hours. While attempting to verify $\text{Helios}^{\text{ped}}$ the tool ran out of memory and we were only able to verify a version of $\text{Helios}^{\text{ped}}$ with two honest voters and no dishonest voter. As the tool is still in a prototype status

we are confident that future optimizations will allow the tool to scale up to this kind of protocols. The tool and example files are available at <https://github.com/ciobaca/akiss>.

ProVerif. The ProVerif tool [4] is an automatic cryptographic protocol verifier. It is based on the representation of protocols by Horn clauses and relies on several approximations. ProVerif can handle several types of properties and in particular equivalence based properties [5] like the privacy-type ones which we are interested in this work. Moreover, ProVerif can handle many different cryptographic primitives, including Pedersen commitments as our case studies show.

The ProVerif tool works by translating biprocesses into Horn clauses built over the two predicates `attacker2` and `message2`. For equivalence checking, biprocess is used to represent the pair of processes for which ProVerif is called to check equivalence. The fact `attacker2( $M, M'$ )` means that the attacker can learn the value M (*resp.* M') from the first (*resp.* second) process encoded by the biprocess. The fact `message2( $M, N, M', N'$ )` means that the message N (*resp.* N') has appeared on the channel M (*resp.* M') while executing the first (*resp.* second) process encoded by the biprocess.

As for the AKISS tool, our extension of ProVerif consists in the addition of constructs for declaring *everlasting channels* and a *future equational theory* (different from the *present* one). We introduce the extra binary predicate `attacker2_ev` for the generation of Horn clauses from biprocesses of our extended ProVerif language. The fact `attacker2_ev( $M, M'$ )` means that in the future, the attacker will either remember or be able to compute from messages he remembers, the value M (*resp.* M'). The declaration of an everlasting channel c generates the following *inheritance* Horn clause:

$$\text{message2} : c[], xm, c[], ym \rightarrow \text{attacker2_ev} : xm, ym$$

This clause transports messages sent on the everlasting channel to the “future”. The declaration of future equations generates the same equations as present ones but using our new `attacker2_ev` predicate. For example, the everlasting equation

$$\text{break}(\text{aenc}(\text{pk}(xk), xr, xm)) = xk$$

will generate the two following clauses

$$\begin{aligned} \text{attacker2_ev} : x, \text{aenc}(\text{pk}(xk), xr, xm) &\rightarrow \text{attacker2_ev} : \text{break}(x), xk \\ \text{attacker2_ev} : \text{aenc}(\text{pk}(xk), xr, xm), x &\rightarrow \text{attacker2_ev} : xk, \text{break}(x) \end{aligned}$$

These clauses model the “future” ability of the attacker to recover the decryption key of ciphertexts he remembers.

Using our extension of the ProVerif tool, we managed to find the attack on $\text{Helios}^{\text{noreplay}}$ presented in section 4.2, but also to prove that $\text{Helios}^{\text{noid}}$, $\text{Helios}^{\text{pedersen}}$ and that Moran – Naor satisfy everlasting vote privacy. However, because of the abstractions made by ProVerif, we had to adapt our models of $\text{Helios}^{\text{noid}}$ and $\text{Helios}^{\text{pedersen}}$ in order for ProVerif to succeed in proving that satisfy everlasting privacy. Indeed, these

	AKISS	ProVerif
Helios	attack on privacy	attack on privacy
Helios ^{noreplay}	proof of privacy attack on everlasting privacy	proof of privacy attack on everlasting privacy
Helios ^{noid}	proof of everlasting privacy	proof of everlasting privacy (voters casting their votes in fixed order)
Helios ^{ped}	proof of everlasting privacy (2 honest voters only)	proof of everlasting privacy (voters casting their votes in fixed order)
Moran-Naor	proof of everlasting privacy	proof of everlasting privacy

Fig. 3. Automated verification using AKISS and ProVerif

two protocols do not satisfy uniformity under reductions, and ProVerif reported false attacks on these two protocols. To overcome this limitation of ProVerif, we fixed the order in which the three voters cast their votes.

The tool and example files are available at <http://markryan.eu/research/EverlastingPrivacy/>.

7 Conclusion

The key idea of “practical” everlasting privacy is that in the future, an attacker will be more powerful in terms of computation (he may be able to break the cryptography) but less powerful in terms of the data he can operate on (transactions between a vote client and the vote server may not have been stored). We realized this idea in the “symbolic” model by allowing different equational theories in different phases, and restricting the information flow from the earlier phase to the later one. We modified ProVerif and AKISS to verify our examples automatically.

We foresee to apply our results to more evolved case studies, e.g. taking into account the zero knowledge proofs that we omitted here for simplicity. Our case studies also show the limitations of the tools for checking equivalence properties which motivates further work to increase their efficiency and scope. Finally, the ability to model different equational theories with restricted information passing between them opens up possibilities for modeling breakable cryptography and other kinds of forward security. In particular it would be interesting to apply the notion of everlasting security to other flavors of anonymity and untraceability.

Acknowledgements. The research leading to these results has received funding from the European Research Council under the European Unions Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 258865, project ProSecure, the ANR projects ProSe (decision ANR 2010-VERS-004) and JCJC VIP (decision ANR-11-JS02-006). We also acknowledge funding from EPSRC projects EP/G02684X/1 “Trustworthy Voting Systems” and EP/H005501/1 “Analysing Security and Privacy Properties”.

References

1. Abadi, M., Fournet, C.: Mobile values, new names, and secure communication. In: 28th Symposium on Principles of Programming Languages, POPL 2001. ACM Press (2001)
2. Adida, B.: Helios: web-based open-audit voting. In: 17th Conference on Security Symposium, SS 2008. USENIX Association (2008)
3. Backes, M., Hritcu, C., Maffei, M.: Automated verification of remote electronic voting protocols in the applied pi-calculus. In: 21st IEEE Computer Security Foundations Symposium, CSF 2008. IEEE (2008)
4. Blanchet, B.: An efficient cryptographic protocol verifier based on Prolog rules. In: 14th Computer Security Foundations Workshop, CSFW 2001. IEEE (2001)
5. Blanchet, B., Abadi, M., Fournet, C.: Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming* 75(1) (2008)
6. Chadha, R., Ciobăcă, Ș., Kremer, S.: Automated Verification of Equivalence Properties of Cryptographic Protocols. In: Seidl, H. (ed.) ESOP 2012. LNCS, vol. 7211, pp. 108–127. Springer, Heidelberg (2012)
7. Chaum, D., Ryan, P.Y.A., Schneider, S.: A Practical Voter-Verifiable Election Scheme. In: De Capitani di Vimercati, S., Syverson, P.F., Gollmann, D. (eds.) ESORICS 2005. LNCS, vol. 3679, pp. 118–139. Springer, Heidelberg (2005)
8. Clarkson, M., Chong, S., Myers, A.: Civitas: Toward a secure voting system. In: 29th IEEE Symposium on Security and Privacy, S&P 2008 (2008)
9. Cortier, V., Smyth, B.: Attacking and fixing helios: An analysis of ballot secrecy. In: 24th IEEE Computer Security Foundations Symposium, CSF 2011. IEEE (June 2011)
10. Cuvelier, E., Peters, T., Pereira, O.: Election verifiability or ballot privacy: Do we need to choose? *SecVote*, Dagstuhl (2012), secvote.uni.lu/slides/opereira-verif-or-priv.pdf
11. Delaune, S., Kremer, S., Ryan, M.D.: Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security* 17(4), 435–487 (2009)
12. Demirel, D., Van De Graaf, J., Araújo, R.: Improving helios with everlasting privacy towards the public. In: International conference on Electronic Voting Technology/Workshop on Trustworthy Elections, EVT/WOTE 2012. USENIX Association (2012)
13. Dreier, J., Lafourcade, P., Lakhnech, Y.: Defining Privacy for Weighted Votes, Single and Multi-voter Coercion. In: Foresti, S., Yung, M., Martinelli, F. (eds.) ESORICS 2012. LNCS, vol. 7459, pp. 451–468. Springer, Heidelberg (2012)
14. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: ACM Workshop on Privacy in the Electronic Society, WPES 2005. ACM (2005)
15. Kremer, S., Ryan, M.: Analysis of an Electronic Voting Protocol in the Applied Pi Calculus. In: Sagiv, M. (ed.) ESOP 2005. LNCS, vol. 3444, pp. 186–200. Springer, Heidelberg (2005)
16. Kremer, S., Ryan, M., Smyth, B.: Election Verifiability in Electronic Voting Protocols. In: Gritzalis, D., Preneel, B., Theoharidou, M. (eds.) ESORICS 2010. LNCS, vol. 6345, pp. 389–404. Springer, Heidelberg (2010)
17. Küsters, R., Truderung, T., Vogt, A.: Accountability: definition and relationship to verifiability. In: ACM Conference on Computer and Communications Security, CCS 2010 (2010)
18. Moran, T., Naor, M.: Receipt-Free Universally-Verifiable Voting with Everlasting Privacy. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 373–392. Springer, Heidelberg (2006)
19. Pedersen, T.P.: Non-interactive and Information-Theoretic Secure Verifiable Secret Sharing. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (1992)