

How Not to Prove Yourself: Pitfalls of the Fiat-Shamir Heuristic and Applications to Helios

David Bernhard¹, Olivier Pereira², and Bogdan Warinschi¹

¹ University of Bristol

² Université Catholique de Louvain

Abstract. The Fiat-Shamir transformation is the most efficient construction of non-interactive zero-knowledge proofs.

This paper is concerned with two variants of the transformation that appear but have not been clearly delineated in existing literature. Both variants start with the prover making a commitment. The strong variant then hashes both the commitment and the statement to be proved, whereas the weak variant hashes only the commitment. This minor change yields dramatically different security guarantees: in situations where malicious provers can select their statements adaptively, the weak Fiat-Shamir transformation yields unsound/unextractable proofs. Yet such settings naturally occur in systems when zero-knowledge proofs are used to enforce honest behavior. We illustrate this point by showing that the use of the weak Fiat-Shamir transformation in the Helios cryptographic voting system leads to several possible security breaches: for some standard types of elections, under plausible circumstances, malicious parties can cause the tallying procedure to run indefinitely and even tamper with the result of the election.

On the positive side, we define a form of adaptive security for zero-knowledge proofs in the random oracle model (essentially simulation-sound extractability), and show that a variant which we call *strong Fiat-Shamir* yields secure non-interactive proofs.

This level of security was assumed in previous works on Helios and our results are then necessary for these analyses to be valid. Additionally, we show that strong proofs in Helios achieve non-malleable encryption and satisfy ballot privacy, improving on previous results that required CCA security.

1 Introduction

Zero-knowledge proofs of knowledge allow a prover to convince a verifier that she holds information satisfying some desirable properties without revealing anything else. To be useful, such proof systems should satisfy completeness (the prover can convince the verifier that a true statement is indeed true) and soundness (the prover cannot convince the verifier that a false statement is true). Zero-knowledge proofs can either be interactive or non-interactive; for the latter

the prover only sends his proof and the verifier decides to accept or reject the statement without any further interaction.

The focus of this paper is on the most common and efficient construction of non-interactive proofs, namely the Fiat-Shamir heuristic [1]. Here, one begins with an interactive sigma protocol, a special type of three-move protocol in which the prover sends a commitment, the verifier answers with a random challenge and the prover completes the protocol with a response. The idea behind the transformation is simple and appealing: have the prover compute the message of the verifier as the hash of the message sent by the prover — if the hash is modelled as a random oracle the message computed this way should look random as in an interactive execution, hence the properties of the original proof system should somehow be preserved.

The transformation appears in the literature in two different forms, depending on what is hashed. In the formalization of Bellare and Rogaway [2], which we refer to as the weak Fiat-Shamir transformation (wFS), the hash takes only the prover's first message as input. Other papers e.g. [3,4] suggest including the statement to be proved in the hash input. In the remainder of the paper we call this the strong Fiat-Shamir transformation (sFS).

CONTRIBUTIONS. The contributions of this paper fall into two main categories. First we identify weaknesses of the weak (sic!) Fiat-Shamir transformation and show that in applications it can be a serious source of insecurity. Secondly, we provide several positive results regarding the strong Fiat-Shamir transformation and its uses in applications.

Insecurity of wFS and Attacks on Helios. Our first results show that the security proofs commonly given for Fiat-Shamir proofs do not hold when applied to weak proofs and when the prover can choose his statement(s) to prove adaptively. This may or may not render a protocol using them insecure, as a protocol may have other means of dealing with adaptivity. For example, in the original application to identification protocols, weak proofs are sufficient.

As an example where weak proofs do not yield security, we consider Helios [5,6], a cryptographic voting protocol that has been used in practice. Versions of Helios have been employed, for example, for the election of the president of the Université catholique de Louvain [6], the Princeton University Undergraduate Student Government [7] and the board of the IACR [8]. We focus on the zero-knowledge protocols implemented since Helios 2.0 [6] for elections based on homomorphic tallying, which are still used in the latest version of Helios as documented on [9] at the time of writing. In brief, those elections work as follows. Trustees first jointly generate an election public key, using NIZK proofs to make sure that this key actually includes contributions from all trustees. Then, to cast a ballot, a voter encrypts a vote and attaches NIZK proofs that the vote is legal. All ballots are placed on a publicly readable bulletin board. Eventually, the election administrators homomorphically add all ballots, decrypt the result and use NIZK to prove the correctness of their actions. The encryption scheme is exponential ElGamal and the particular NIZKs involved are obtained by applying the weak Fiat-Shamir transformation to the Schnorr [10], Chaum-Pedersen [11]

and disjunctive Chaum-Pedersen protocols (and variants thereof). These proofs are used to guarantee that the privacy of the votes rests on all trustees, to enforce that voters create ballots containing valid votes and to prevent dishonest administrators from claiming a wrong result. We show that the use of the wFS transformation is the source of three types of insecurity:

- a) breaking verifiability by allowing colluding administrators to cast a single ballot that is not well-formed and contains any chosen number of votes for a specific candidate,
- b) breaking liveness of the system by allowing colluding administrators to fail providing the election outcome while proving that they behave honestly, or by allowing voters to cast a random vote which leads to tallying taking superpolynomial time, and
- c) breaking privacy by allowing the casting of related ballots that do not contain mere copies of previously submitted ciphertexts.

The first two of these attacks are undetectable under normal circumstances.

While our focus is on Helios which is our motivating application, in the full version of our paper we also show attacks against schemes constructed via the Naor-Yung paradigm and via the encrypt-then-prove construction: when using proofs derived through wFS these constructions may yield malleable encryption schemes.

Security of Strong Fiat-Shamir and Applications. The problems that we have identified in the use of the wFS do not apply to proofs obtained through the strong version of the transformation. It is then natural to ask what level of security does one get from these proofs. We provide several results. First, we formulate a security notion for non-interactive zero-knowledge proofs of knowledge which captures adversaries that can choose their statements adaptively. In essence, this notion is the analogue of simulation-sound extractability defined by Groth in the common reference string model [12]. Informally, a malicious prover is allowed to see simulated proofs (of potentially fake statements) and aims to provide valid looking proofs for adaptively chosen statements in such a way that an extractor cannot obtain witnesses. Interestingly, our definition is not simply a rehashing of the notion in [12]. In the random-oracle model, extraction requires the rewinding of the prover (as opposed to merely using a trapdoor) and in turn, this implies complex interaction between the adversary, the simulator and the extractor. We then show that applying sFS to Σ -protocols results in protocols that are simulation-sound extractable. Our result seems to be the first thorough investigation on the precise security guarantees offered by such proofs.

As a first application of this result, we investigate the security of non-malleable encryption schemes that are built by combining an IND-CPA encryption scheme with a proof of knowledge of the randomness used in the encryption process. We refer to this construction as the $Enc+PoK$ approach. A well-known instantiation is the TDH0 scheme introduced and studied by Shoup and Gennaro [13]. Intuitively the construction should achieve IND-CCA security but so far, all attempts have failed to confirm or disprove this under natural assumptions (e.g., DDH in the random oracle model) [14,15]. As a consequence, the form of non-malleability ensured by $Enc+PoK$ schemes is, surprisingly, still unknown. We

provide a lower-bound on the answer to this question: if the proof of knowledge used in the encryption process is simulation-sound extractable, then the resulting scheme is NM-CPA secure. An immediate corollary is that the TDH0 scheme is NM-CPA secure in the random oracle model under the DDH assumption.

We then turn to the analysis of ballot privacy in Helios. Prior work shows that ballot privacy is guaranteed if the encryption scheme used in the construction is IND-CCA [16,17]. Since ballots in Helios use the Enc+PoK paradigm (which, as discussed above, is not known to be IND-CCA) a natural suggestion is then to replace it with something stronger. For example, Bernhard et al. suggested applying the Naor-Yung transformation to the underlying ElGamal encryption [17], while Bulens et al. used a variant of the TDH2 scheme [18]. These modifications both substantially increase the computational costs of the system and require major changes in the implementation.

Our final result is to show that although the NM-CPA notion is strictly weaker than IND-CCA [19], it is sufficient to ensure ballot privacy. In particular a minor tweak of the Enc+PoK construction currently used in Helios where we replace wFS with its strong counterpart and check for repeated ciphertexts is sufficient. The change that we require is easily accomplished by including additional elements in the inputs of the hash function and preserves the current level of efficiency.

2 The Fiat-Shamir/Blum Transformation

In this section we introduce the two variants of the Fiat-Shamir heuristic that we analyze. We start by fixing notation and recalling some standard notions. In the following we let $R \subseteq \mathcal{P}(\{0, 1\}^* \times \{0, 1\}^*)$ be an efficiently computable relation. R defines a language $\mathcal{L}_R = \{Y \in \{0, 1\}^* \mid \exists w : R(w, Y)\}$ in NP. We further assume that there is a well-defined set $A \supseteq \mathcal{L}$ decidable in polynomial time.¹

A non-interactive proof system for language \mathcal{L}_R is a pair algorithms (Prove, Verify). Such a proof system is complete for \mathcal{L}_R if for every $(w, Y) \in R$, with overwhelming probability if $\pi \leftarrow \text{Prove}(w, Y)$ then $\text{Verify}(Y, \pi) = 1$. We define soundness of such proof systems (the property that a cheating prover cannot make the verifier accept a false statement) later in the paper. Here we recall the notion of zero-knowledge in the random oracle model [20].

In this setting, a simulator \mathcal{S} for a proof system is an algorithm in charge of answering random oracle queries and producing valid proofs for any statement $Y \in A$ with respect to this oracle. In particular, it can “patch” the oracle to create its simulated proofs. Such a simulator responds to the following queries:

¹ Suppose that \mathcal{L} is the set of DDH triples (G^a, G^b, G^{ab}) over some group \mathbb{G} . Then A could be \mathbb{G}^3 . The reason for defining this formally is that we will later expect our zero-knowledge simulator to produce valid “proofs” for some “false statements”, but which ones? Can it produce a proof for the statement consisting of the empty string, for example? We use A as the class of statements on which the simulator can produce proofs.

$\mathcal{H}(s)$ \mathcal{S} maintains a list of oracle query/response pairs. For repeated queries, \mathcal{S} answers consistently; for fresh queries, \mathcal{S} draws a random value r , adds (s, r) to its list and returns r .

Simulate(Y) For $Y \in \Lambda$, the simulator returns a proof π such that $\text{Verify}(Y, \pi) = 1$ if the verifier uses the simulator for its oracle queries. \mathcal{S} can add query/response pairs to its oracle list to process a simulation query.

Definition 1 (Zero-Knowledge). *A proof system is zero-knowledge if there is a simulator \mathcal{S} such that no adversary who can make queries to the random oracle and queries of the form $\text{create-proof}(w, Y)$ can distinguish the following two settings with non-negligibly better than $1/2$ probability.*

1. Random oracle queries are answered by a random oracle. In response to $\text{create-proof}(w, Y)$, the challenger checks that $R(w, Y)$. If not, he returns \perp . Otherwise, he returns $\text{Prove}(w, Y)$.
2. The challenger runs a copy of the simulator \mathcal{S} . It forwards random oracle queries to \mathcal{S} directly. For $\text{create-proof}(w, Y)$, the challenger checks if $R(w, Y)$ holds: if not, the challenger returns \perp ; if it holds, the challenger sends $\text{Simulate}(Y)$ to \mathcal{S} and returns the result to the adversary.

Sigma Protocols. A sigma protocol for a language \mathcal{L}_R is a protocol for two parties, a prover and a verifier. Both share a statement $Y \in \mathcal{L}_R$ as input and the prover may additionally hold a witness w .

The prover begins by sending a value A known as the commitment. The verifier replies with a challenge c drawn uniformly from a fixed challenge set. The prover finishes the protocol with a response f whereupon the verifier applies a deterministic algorithm Verify to Y, A, c and f which can accept or reject this execution.

A sigma protocol is correct (w.r.t. \mathcal{L}_R) if the prover, on input a pair (w, Y) satisfying R and interacting with the verifier who has input Y , gets the verifier to accept with probability 1.

A sigma protocol has *special honest verifier zero knowledge* if there is an algorithm Simulate that takes as input a statement $Y \in \Lambda$, challenge c and response f and outputs a commitment A such that $\text{Verify}(Y, A, c, f) = 1$ and furthermore, if c and f were chosen uniformly at random from their respective domains then the triple (A, c, f) is distributed identically to that of an execution between the prover and the verifier. Notice that the verifier is supposed to work with statements that may be false.

A sigma protocol has *special soundness* if there is an algorithm Extract that takes as input a statement Y and any two triples (A, c, f) and (A', c', f') such that both verify w.r.t. Y , $A = A'$ and $c \neq c'$, and returns a witness w such that $R(w, Y)$.

The Fiat-Shamir Transformation. The Fiat-Shamir transformation [1] (which [2] attributes to Blum) is a technique to make sigma protocols non-interactive using a cryptographic hash function. There are two commonly used descriptions of this technique that we call weak and strong Fiat-Shamir and which we describe together in the following definition.

Definition 2 (Fiat-Shamir Transformation). Let $\Sigma = (\text{Prove}_\Sigma, \text{Verify}_\Sigma)$ be a sigma protocol and \mathcal{H} a hash function. The weak Fiat-Shamir transformation of Σ is the proof system $\text{wFS}_{\mathcal{H}}(\Sigma) = (\text{Prove}, \text{Verify})$ defined as follows:

Prove(w, Y) Run $\text{Prove}_\Sigma(w, Y)$ to obtain commitment A . Compute $c \leftarrow \mathcal{H}(A)$. Complete the run of Prover_Σ with c as input to get the response f . Output the pair (c, f) .

Verify(Y, c, f) Compute A from (Y, c, f) , then run $\text{Verify}_\Sigma(Y, A, c, f)$.

The strong Fiat-Shamir transformation of Σ , i.e., $\text{sFS}(\Sigma) = (\text{Prove}, \text{Verify})$ is obtained as above with the difference that c is computed by $c \leftarrow \mathcal{H}(Y, A)$.

3 Pitfalls of the Weak Fiat-Shamir Transformation

We now describe various standard protocols in which the use of the weak Fiat-Shamir transformation can have undesirable effects. We illustrate these effects through several new practical attacks on various components of the Helios voting system, which relies on these protocols.

SCHNORR PROOFS. The Schnorr [10] signature scheme is the weak Fiat-Shamir transformation of the Schnorr identification protocol. In a group \mathbb{G} of order q generated by G , it proves knowledge of an exponent x satisfying the equation $X = G^x$ for a known X . Viewing (x, X) as a signing/verification key pair and including a message in the hash input yields a signature of knowledge.

To create a proof, the prover picks a random $a \leftarrow \mathbb{Z}_q$ and computes $A = G^a$. He then hashes A to create a challenge $c = \mathcal{H}(A)$. Finally he computes $f = a + cx$; the proof is the pair (c, f) and the verification procedure consists in checking the equation $c \stackrel{?}{=} \mathcal{H}\left(\frac{G^f}{X^c}\right)$.

The weak Fiat-Shamir transformation can safely be used here, as discussed in previous analysis [10,21], since the public key X is selected first and given as input to the adversary who tries to produce a forgery.

However, if the goal of the adversary is to build a valid triple (X, c, f) for any X of his choice, then this protocol is not a proof of knowledge anymore unless the discrete logarithm problem is easy in \mathbb{G} . Suppose indeed that there is an extractor \mathcal{K} that, by interacting with any prover \mathcal{P} that provides a valid triple (X, c, f) , extracts $x = \log_G(X)$. This extractor can be used to solve an instance Y of the discrete logarithm problem with respect to (\mathbb{G}, G) as follows: use Y as the proof commitment, compute $c = \mathcal{H}(Y)$, choose $f \leftarrow \mathbb{Z}_q$ and set $X = \left(\frac{G^f}{Y}\right)^{\frac{1}{c}}$. Since the proof (Y, c, f) passes the verification procedure for statement X , the extractor \mathcal{K} should be able to compute $x = \log_G(X)$ by interacting with our prover. We now observe that, by taking the discrete logarithm in base G on both sides of the definition of X , we obtain the solution $\log_G(Y) = f - cx$ to the discrete logarithm challenge.

Application to Helios. Schnorr proofs are used during the key generation procedure of Helios as a way to prevent trustees from choosing their public key as a function of the public key of the other trustees, which could give them the

possibility to select the election private key at will and to decrypt all individual votes [22]. While the scenario above shows that trustees who publish a public key together with a Schnorr proof for that public key do not necessarily know the corresponding private key, the fact that our scenario does not allow the prover to choose his statement (but just to compute it as a function of the elements of the proof) does not seem to give rise to any practical attack. These weak Schnorr proofs would, however, break the proof of ballot privacy that we give later in this paper (assuming strong proofs).

CHAUM-PEDERSEN PROOFS. Chaum and Pedersen [11] introduced a proof of discrete logarithm equality, which they make non-interactive using the strong form of the Fiat-Shamir transformation. More precisely, given two group elements (G, X) , a prover who knows the discrete logarithm $x = \log_G(X)$ can prove that two group elements (R, S) satisfy the relation $\log_G(X) = \log_R(S)$ as follows. He picks a random $a \leftarrow \mathbb{Z}_q$, computes $A = G^a$, $B = R^a$, $c = \mathcal{H}(R, S, A, B)$ and $f = a + cx$. The proof is the pair (c, f) and the verification procedure consists in checking the equation $c \stackrel{?}{=} \mathcal{H}(R, S, \frac{G^f}{X^c}, \frac{R^f}{S^c})$.

We observe that this proof is not sound anymore if it is used as a proof that three elements (X, R, S) are such that $\log_G(X) = \log_R(S)$, that is, if the prover also has the possibility to choose X in the process of building his proof. Indeed, a prover could select $(a, b, r, s) \leftarrow \mathbb{Z}_q^4$ at random, compute $A = G^a$, $B = G^b$, $R = G^r$ and $S = G^s$ from which he can compute $c = \mathcal{H}(R, S, A, B)$ and $f = \frac{b+cs}{r}$. He now completes the proof by computing $x = (f - a)/c$ and setting $X = G^x$. Now, we observe that $\log_G(X) = \frac{s}{r} + \frac{b-ar}{rc}$ while $\log_R(S) = \frac{s}{r}$, which differ with overwhelming probability.

Application to Helios. Chaum-Pedersen proofs instantiated with the weak Fiat-Shamir transformation (that is, $c = \mathcal{H}(A, B)$) are used during the ElGamal decryption procedure of Helios, in order to demonstrate that the decryption of the product of the votes that is computed by the trustees is consistent with the public key. More precisely, given a public key X and a ciphertext (R, S) that encrypts the sum of all votes, a trustee is required to compute $T = R^x$ where $x = \log_G(X)$ and to publish it together with a Chaum-Pedersen proof that $\log_G(X) = \log_R(S)$. The ElGamal decryption is then computed as $\log_G(S/T)$.

In this proof, a malicious trustee does not have the possibility to choose his private key at decryption time, but has the possibility to select T as part of the proof computation process. He can do so as follows. Select $(a, b) \leftarrow \mathbb{Z}_q^2$ at random, compute the proof commitments $A = G^a$ and $B = G^b$, the challenge $c = \mathcal{H}(A, B)$ and the response $f = a + cx$. Eventually, compute the decryption factor $T = (\frac{R^f}{B})^{\frac{1}{c}}$. It is easy to verify that the proof (c, f) is valid for the tuple (G, X, R, S) , but that $\log_R(T) = x + \frac{ar-b}{c}$, which will be different from x with overwhelming probability. As a result, the decryption procedure will provide an aberrant result: an essentially random element of \mathbb{Z}_q . This strategy provides a way to build a denial of service attack against a Helios election, without anyone being able to detect who was responsible.

A more dangerous attack can be mounted if we assume that the trustees have the possibility to passively eavesdrop on the randomness of all voters. Though demanding, such an attack is still easier to mount and harder to detect than a full active attack. We would expect the impact of such a scenario to be “only” a complete loss of privacy by the voters but we show that it actually provides a way for the trustees to announce any election outcome of their choice as soon as they can actively corrupt a single voter (which can happen simply if a trustee is a voter himself).

Consider an election with trustees who would like to announce the election outcome m . These trustees select a private key x and publish the public key X . They also select $(a, b) \leftarrow \mathbb{Z}_q^2$, compute $A = G^a$, $B = G^b$, $c = H(A, B)$ and $f = a + cx$. Then all the voters submit their votes, except the corrupted one who waits until the last minute of the election. At that time, the trustees compute the product of all encrypted votes that have been submitted and obtain a ciphertext $(R', S') = (G^{r'}, G^{m'} \cdot G^{xr'})$ for some values r' and m' that they can compute using the randomness of the voters. They now compute $r = \frac{b+c(m'-m)}{f-cx}$, as well as a ciphertext $(G^{r-r'}, G^{x(r-r')})$ which is an encryption of 0 for which they can compute a proof of validity since they know $r - r'$. This ciphertext and proof are submitted by the corrupted voter, with the effect that the product of all encrypted votes is $(R, S) = (G^r, G^{m'} \cdot G^{xr})$. It can now be verified that (c, f) form a valid proof that $\log_G(X) = \log_R(\frac{S}{G^m})$, which indicates that m is the outcome of the election.

DISJUNCTIVE CHAUM-PEDERSEN PROOFS. Disjunctive proofs allow proving that one of two statements holds without revealing which one is correct. These proofs have numerous applications. For instance, they can be used by a voter to demonstrate that a ciphertext he produced is an encryption of either 0 or 1 (but nothing else), expressing whether or not he supports a candidate.

Suppose that a voter builds an exponential ElGamal ciphertext (R, S) with respect to public key X and wants to prove that it encrypts 0 or 1. We consider the case where it is an encryption of 1 (the other case is similar). First, the voter simulates a proof that $\log_R(S) = x$ by selecting a random proof $(c_0, f_0) \leftarrow \mathbb{Z}_q^2$ and computing $A_0 = G^{f_0}/R^{c_0}$ and $B_0 = X^{f_0}/S^{c_0}$. Then he selects $a_1 \leftarrow \mathbb{Z}_q$, computes $A_1 = G^{a_1}$, $B_1 = X^{a_1}$, $c = \mathcal{H}(A_0, B_0, A_1, B_1)$, $c_1 = c - c_0$ and $f_1 = a_1 + c_1 r$. The proof consists of (c_0, c_1, f_0, f_1) and verification consists of verifying whether $c_0 + c_1 = \mathcal{H}(\frac{G^{f_0}}{R^{c_0}}, \frac{X^{f_0}}{S^{c_0}}, \frac{G^{f_1}}{R^{c_1}}, \frac{X^{f_1}}{(S/G)^{c_1}})$.

Application to Helios. The proof we just described is exactly the one used in Helios to guarantee that voters encode at most one vote for each candidate and it exhibits weaknesses that are similar to those described above, but with an even more dangerous effect. Consider an election organized by corrupted trustees who would like to influence the election outcome by adding (or removing) m approvals to a candidate of their choice. These trustees now have the freedom to choose any public key and ciphertext of their choice that would allow them to compute an encryption of m and to prove that it is an encryption of 0 or 1. They can achieve this as follows.

They first select $(a_0, b_0, a_1, b_1) \leftarrow \mathbb{Z}_q^4$, from which they compute the commitments $A_0 = G^{a_0}$, $B_0 = G^{b_0}$, $A_1 = G^{a_1}$, $B_1 = G^{b_1}$, the challenge $c = \mathcal{H}(A_0, B_0, A_1, B_1)$ and the private key $x = \frac{(b_0 + cm)(1-m) - b_1 m}{a_0(1-m) - a_1 m}$ and the public key $X = G^x$. Using this public key, they select a random encryption of m by selecting $r \leftarrow \mathbb{Z}_q$ and computing $(R, S) = (G^r, G^m X^r)$. Eventually, they compute the challenges $c_1 = \frac{b_1 - a_1 x}{1-m}$ and $c_0 = c - c_1$ and the responses $f_0 = a_0 + c_0 r$ and $f_1 = a_1 + c_1 r$. It can be verified that (c_0, c_1, f_0, f_1) form a proof that (R, S) encrypt 0 or 1, while it actually encrypts an arbitrary m .

Furthermore, it can be observed that this proof, like the others that we presented above, is indistinguishable from a regular one.

Other attack possibilities exist, based on the same techniques. For instance, a voter who does not know the election private key can build a ciphertext that encrypts a random value in \mathbb{Z}_q and prove that it encrypts 0 or 1, which would make the decryption procedure fail. We do not know however whether it is possible to build such a proof in a way that is indistinguishable from a regular one.²

ENCRYPT + PoK. Adding a proof of knowledge of the plaintext/randomness to a ciphertext in an IND-CPA secure public key encryption scheme is a common way to yield a non-malleable encryption scheme.³ We formalise this construction and show that using wFS does *not* yield non-malleable encryption.

Definition 3 (Encrypt+PoK). Let $\mathfrak{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a public-key encryption scheme. Let $R((m, r), (Y, pk)) := (Y = \text{Enc}(pk, m; r))$ be the relation that Y is an encryption of m with randomness r for public key pk , let Λ be the ciphertext space (or some suitable superset thereof) and let $\mathfrak{P} = (\text{Prove}, \text{Verify})$ be a NIZK-PoK for this relation.

The *Encrypt+PoK* transformation $\mathfrak{E}_{\mathfrak{P}}$ is the following encryption scheme.

KeyGen' Run KeyGen.

Enc'(pk, m) Draw some randomness r and create a ciphertext $E = \text{Enc}(pk, m; r)$.

Create a proof $\pi \leftarrow \text{Prove}(pk, E, m, r)$. The ciphertext is the pair (E, π) .

Dec'(sk, E, π) First run $\text{Verify}(pk, E, \pi)$. If this fails (returns 0), output \perp and halt. Otherwise, return $\text{Dec}(sk, E)$.

Consider the ElGamal encryption scheme with weak Schnorr proofs of the randomness used for encryption (which would allow one to extract the message), which would be a weak variant of the TDH0 scheme [13]. In other words, a ciphertext for a message M under public key X is $(G^r, M \cdot X^r, c, f)$ where $c = \mathcal{H}(G^f / (G^r)^c)$. We can rerandomise such a ciphertext (R, S, c, f) by picking a random u and setting the new ciphertext to be $(R \cdot G^u, S \cdot X^u, c, f + cu)$. The new plaintext is the same as the old one as $S/R^x = M \cdot X^{r+u} / (G^{r+u})^x = M$ and the proof still verifies as $c = \mathcal{H}\left(\frac{G^{f+cu}}{(G^{r+u})^c}\right) = \mathcal{H}\left(\frac{G^f}{(G^r)^c}\right)$. Clearly, this encryption scheme is malleable.

² Our current technique involves setting $c_0 = 0$. While such a ballot passes the current Helios verifier, this could be detected in an audit.

³ Though the exact form of non-malleability that is provided is unclear [13].

Application to Helios. The same rerandomisation technique can be applied to current Helios ballots, giving a ballot privacy attack in the style of Cortier and Smyth [23] (based on the same principles as the attacks described in [24,25].) Helios ballots contain ElGamal ciphertexts (R, S) with disjunctive Chaum-Pedersen proofs (c_0, c_1, f_0, f_1) . To rerandomise such a ciphertext, pick a random u and set $R' = R \cdot G^u$, $S' = S \cdot Y^u$, $f'_0 = f_0 + c_0u$ and $f'_1 = f_1 + c_1u$. Unlike previously known rerandomisation techniques, this one does not make use of a repeated ElGamal ciphertext or proof. It can be detected however by checking for repeated hash values, just as for the previous attacks.

FURTHER EXAMPLES. The various attacks that we described above focus on applications to the Helios voting system, which uses the weak Fiat-Shamir transformation in all proofs. We believe that these examples provide clear evidence that the weak Fiat-Shamir transformation should not be used in that context: in particular, we showed that malicious authorities can arbitrarily influence the outcome of an election, which is in clear contradiction with the universal verifiability properties expected from that system. In the next sections, we will focus on the properties of the strong Fiat-Shamir transformation and show the benefits that its adoption would provide for the Helios system.

We stress that there are various other contexts in which the weak Fiat-Shamir transformation should not be used. For instance, similarly to our observation for the weak variant of the TDH0 scheme, the scheme resulting from the Naor-Yung transformation [26] applied to ElGamal encryption may become malleable if the weak Fiat-Shamir transformation is used, contradicting the level of desired security. We provide attacks against a concrete instantiation of that transformation in the full version of our paper.

4 Simulation Sound Extractable Proofs

The examples discussed in the previous section show that the wFS transform fails to offer even the most basic soundness properties in many contexts. We now investigate the soundness properties of the sFS transform. More precisely, we formulate the notion of *simulation sound extractable proofs* in the random oracle model and show its applications to the sFS transformation. Our definition draws inspiration from that of witness-extended emulation [4] in which the existence of an extractor is demanded such that for any adversary returning a vector of statements and proofs, the extractor returns identically distributed elements along with the witnesses to the proven statements. However, the definition is perhaps more appropriately viewed as the analogue definition of simulation sound extractability defined by Groth [12] which combines the simulation soundness approach of Sahai [27], with proofs of knowledge [28].

We consider a malicious prover who may ask to see simulated proofs (as in simulation-soundness). The extractor that we consider gets the transcript of a run of the prover where the prover outputs several valid proofs together with the transcript of random oracle queries. His goal is to extract witnesses of these proofs. In the process, we allow the extractor to invoke and communicate with

copies of the prover that use the same randomness as the run it is trying to extract from. This ability is what permits the knowledge extractor to fork the prover's execution without giving the extractor access to the coins of the prover.⁴ A bit more precisely, a malicious prover for a proof system $\mathfrak{P} = (\text{Prove}, \text{Verify})$ is an algorithm \mathcal{A} that expects access to two oracles: a hashing oracle and a simulation oracle. Thus \mathcal{A} may submit some string s and expects $\mathcal{H}(s)$ in return and it may also make simulation calls $\text{Simulate}(Y)$ for any statement $Y \in \Lambda$ and expects to obtain a proof π such that $\text{Verify}(Y, \pi) = 1$. The prover returns a pair of vectors $(\mathbf{Y}, \boldsymbol{\pi})$.

Definition 4 (Simulation Sound Extractability). *Let \mathfrak{P} be a zero-knowledge proof system with simulator \mathcal{S} . We say that \mathfrak{P} is simulation sound extractable (SSE) if there exists an extractor \mathcal{K} such that for every prover \mathcal{A} , \mathcal{K} wins the following game with non-negligible probability.*

1. (*Initial run.*) The game selects a random string ω for \mathcal{A} . It runs an instance of \mathcal{A} with the simulator \mathcal{S} until \mathcal{A} makes his output and halts. If \mathcal{A} does not output any proofs, any of the proofs do not verify (w.r.t. the instance of \mathcal{S} used as the random oracle) or any of \mathcal{A} 's statement/proof pairs (Y, π) is such that π was the result of a $\text{Simulate}(Y)$ query, then \mathcal{K} wins the game directly.
2. (*Extraction.*) The game runs an instance of \mathcal{K} , giving it the transcript of all queries in the initial run and the produced $(\mathbf{Y}, \boldsymbol{\pi})$ as input. \mathcal{K} may repeatedly make one type of query invoke in response to which the game runs a new invocation of \mathcal{A} on the same randomness ω that it chose for the initial run. All queries made by these instances are forwarded to \mathcal{K} who can reply to them.
3. \mathcal{K} wins the game if it can output a vector of witnesses \mathbf{w} that match the statements \mathbf{Y} of the initial run, i.e. for all i we have $R(w_i, Y_i)$.

The following theorem confirms that the strong Fiat-Shamir transformation yields proof systems that satisfy the notion we described above.

Theorem 1. *Let Σ be a sigma protocol with a challenge space that is exponentially large in the security parameter, special soundness and special honest verifier zero-knowledge. Then $\text{sFS}(\Sigma)$ is zero-knowledge and simulation sound extractable with respect to expected polynomial-time adversaries.*

Applications. The Schnorr and Chaum-Pedersen protocols are clearly both sigma protocols with special soundness and special honest verifier zero knowledge so Theorem 1 applies and the sFS versions of these protocols are SSE proofs. For disjunctive Chaum-Pedersen, the challenge is the actual c obtained from the verifier and the response is the tuple $f = (f_0, f_1)$. This is a sigma protocol with special soundness and almost special honest verifier zero knowledge — almost, because in our definition the simulator chooses c, f independently and uniformly

⁴ Although not necessary for this paper, hiding the adversary's randomness from the extractor can be helpful in other contexts to prove separation results.

at random yet if $c \neq c_0 + c_1$ then the resulting proof will not verify, patched oracle or not. We could fix this by not sending c_1 in the response and having the verifier recompute $c_1 = c - c_0$. We will ignore this point as it is easy to see that, if the simulator chooses c, f at random and then adjusts c_0 , all relevant theorems still hold. In particular, the sFS transformation of disjunctive Chaum-Pedersen is still a simulation-sound extractable proof.

Encrypt + PoK.

With this notion we can restore the folklore result that appending a PoK to an IND-CPA scheme gives a NM-CPA one, if the PoK is simulation-sound extractable. For space reasons the proof is only in the full version of our paper.

Theorem 2. *Let \mathfrak{E} be an IND-CPA secure encryption scheme and \mathfrak{P} be a simulation-sound extractable NIZK-PoK for the encryption relation. Then $\mathfrak{E}_{\mathfrak{P}}$ is non-malleable (NM-CPA) secure with respect to expected polynomial-time adversaries.*

5 Ballot Privacy in Helios

In this section we propose a modification to Helios and prove that it satisfies ballot privacy in the model of single-pass voting of Bernhard et al. [17].

SINGLE-PASS SCHEMES. A single-pass voting scheme is a protocol consisting of the following algorithms and execution protocol for a set \mathcal{V} of voters, a set \mathcal{T} of trustees and a bulletin board \mathcal{B} . The class of single-pass schemes includes not only Helios [5] but also several other cryptographic voting schemes [29,30,27,31]. Single-pass voting models two of the most popular approaches to cryptographic voting, homomorphic tallying and mix-nets. Voters need only read a single message off the board (the election specification and public keys) and post a single message (their ballot) in return. We assume some underlying voter authentication mechanism.⁵

Setup(1^λ) is an algorithm to create public parameters for the election and secret ones for all trustees.

Setup produces one public output Y known as the public key of the election and a secret output x_i for each trustee \mathcal{T}_i in the set of trustees \mathcal{T} . The secret outputs of all trustees together are known as the secret key of the election.

Vote(id, v, Y) is a probabilistic algorithm run by voters to prepare their votes for submission. It takes as input a voter's identity id , a vote v and public information Y and outputs a ballot $s \leftarrow \text{Vote}(id, v, Y)$.

Validate(b, s) models the algorithm run by the bulletin board during voting. Its inputs are the current board state b and the submitted ballot s . It returns 1 if the submission is deemed valid (given the current state of the board) and 0 otherwise.

⁵ In the election of the president of UC Louvain [6] using Helios, authentication was handled by the university's existing infrastructure. As such it escapes cryptographic modelling and we choose not to model authentication (in particular we do not wish to assume a PKI).

$\text{Tally}(b)$ is a tallying protocol that is run by the trustees. Its inputs are the board so far and the private data kept by the trustees from the setup phase.

$\text{Result}(b)$ is a deterministic algorithm that takes a bulletin board b of a completed election and returns the result of the election, or a special symbol \perp if the board does not contain a valid result.

A single-pass scheme is executed as follows.

1. *Setup phase.* The trustees run the **Setup** algorithm and post the public key Y to the bulletin board.
2. *Voting phase.* Each voter may proceed as follows or abstain. He reads public key Y off the board and computes a ballot $s \leftarrow \text{Vote}(id, v, Y)$ where id is his identity and v is his vote, and submits s to the board.
The board runs $\text{Validate}(b, s)$ on every submission it receives and appends valid ones to its state.
3. *Tallying phase.* The trustees run the **Tally** protocol and may post to the board.

A single-pass protocol is correct w.r.t. a result function ρ if as long as everyone follows the protocol, with overwhelming probability (in the security parameter) none of the algorithms abort, **Result** returns a result when executed on the board at the end of the tallying phase and this result corresponds to ρ evaluated on the votes cast by the voters.⁶

BALLOT PRIVACY. We base our definition of ballot privacy on previous work in this area by Bernhard et al. [17,32]. Ballot privacy is defined by means of a cryptographic indistinguishability game. The new feature of our definition is that it can deal with dishonest trustees; we introduce a simulator to handle tallying in this case.

Definition 5 (Ballot Privacy). *A single-pass protocol for n trustees and any number of voters has ballot privacy against up to $m < n$ dishonest trustees if there is a simulator \mathcal{S} such that for any efficient adversary \mathcal{A} , the advantage $\Pr[\mathcal{A} \text{ wins}] - 1/2$ against the following indistinguishability game is negligible (as a function of the security parameter). The simulator \mathcal{S} is given black-box access to the adversary \mathcal{A} and may invoke further copies of \mathcal{A} using the same randomness as was used in the main run in the security game. We assume static corruption of trustees: the sets of honest and dishonest trustees are fixed in advance. The adversary can adaptively choose voters to be honest or dishonest, however.*

Setup Phase. The challenger picks a bit $\beta \leftarrow \{0, 1\}$ uniformly at random. He sets up two bulletin boards \mathcal{L} and \mathcal{R} . The adversary is given access to either \mathcal{L} if $\beta = 0$ or to \mathcal{R} if $\beta = 1$.

The trustees jointly run the **Setup** protocol, the challenger playing the honest trustees and the adversary, the corrupt ones. This produces some output Y on the visible board. The challenger then copies Y from the visible board to the hidden one. If the setup phase fails to complete, the adversary loses the game.

⁶ One may also want ρ to operate on (v, id) pairs: in the Helios election at UC Louvain, votes from students, faculty and staff were weighted differently.

Voting Phase. The adversary may make two types of queries.

Vote($id, v_{\mathcal{L}}, v_{\mathcal{R}}$) Queries. The adversary provides a voter identity id and two votes $(v_{\mathcal{L}}, v_{\mathcal{R}})$. The challenger runs $b_{\mathcal{L}} \leftarrow \text{Vote}(id, v_{\mathcal{L}}, Y)$ and $b_{\mathcal{R}} \leftarrow \text{Vote}(id, v_{\mathcal{R}}, Y)$, where Y is the public key of the election that can be computed using `Keys` on the public information on the board from the setup phase.

The ballots $b_{\mathcal{L}}$ and $b_{\mathcal{R}}$ are submitted to the corresponding boards which process them normally (run `Validate` and append the ballot if it passes validation).

Ballot(id, b) Queries. These are queries made on behalf of corrupt voters. Here the adversary provides a ballot b . The challenger first submits b to the board visible to the adversary, which validates it and appends it if validation is successful. If the ballot successfully validates on the visible board, the challenger also submits the ballot to the invisible board which again validates the ballot and appends it if successful.

Tallying Phase. If the adversary sees the \mathcal{L} board ($\beta = 0$) then tallying can take place as normal. The trustees execute the Tally protocol, the challenger playing the honest ones and the adversary, the dishonest ones.

If the adversary sees the \mathcal{R} board, the challenger starts up the simulator \mathcal{S} and passes it both the \mathcal{L} and \mathcal{R} boards and the state of the honest trustees. In the random oracle model, the simulator is responsible for the random oracle from this point onwards (and gets a list of all previously queried input/output pairs). The simulator acts on behalf of the honest trustees from now onwards and may post to the board.

At the end of the game the adversary may make a guess of β and wins if his guess is correct.

We propose to fix Helios by changing all proofs to their strong counterparts. This allows us to state the following theorem. The proof along with a detailed description of modified Helios in the single-pass model can be found in the full version of our paper.

Theorem 3. *In the random oracle model, the modified Helios (using strong proofs) satisfies ballot privacy against up to $m = n - 1$ dishonest trustees, assuming that DDH is hard in the underlying group.*

6 Conclusion

The prominence of Helios (it has been used in several real elections, notably in the election of the IACR board of directors) justifies the level of attention it has recently received. Results are divided between finding attacks against ballot privacy (e.g. the method of casting related ballots [23,33] which we further refine in this paper) and proposing modifications that enable rigorous security proofs [23,17,34]. Our paper seems to be the natural convergence point. We identify the use of weak Fiat-Shamir proofs as a source of attacks much stronger than all those previously proposed: we have presented new and unforeseen consequences of

these weak proofs and we have shown that switching to their strong counterpart allows for a proof of ballot secrecy for Helios, and provides a crucial assumption on which existing verifiability analyses of Helios rely [34]. In the process, we have made several conceptual contributions: we have defined simulation sound extractability in the random oracle model, proved that the strong Fiat-Shamir transformation yields secure non-interactive zero-knowledge proofs of knowledge, and justified the new notion through applications that include the Enc+PoK paradigm.

In the remainder of this section we discuss two points that naturally arise from our work.

Usability of wFS. Our results discourage the use of wFS proofs as they may lead to failures in the systems that employ them. Nonetheless, the transformation works well for its original application (and its generalizations [35]) in constructing signature schemes from identification schemes [1], since the statement (essentially the verification key) is fixed in advance. It is interesting to find other settings where wFS can actually be used safely. An intriguing possibility is to exploit malleability of wFS proofs as, for example, in the recent work of Chase et al. [36] that relies on controlled malleability of (standard model) non-interactive zero-knowledge proofs. A necessary first step in this direction is understanding precisely what is the level of malleability of wFS proofs, which we leave for further work.

Practical impact of our attacks. As Helios in its current form has been used in real elections, a discussion of the impact of our attacks in practice is in order. We note that our attacks have been tested and succeeded on the current version of the Helios system on <http://vote.heliosvoting.org>.

Our denial of service attacks may only have an impact on future elections: as far as we know, all Helios elections led to the successful computation of a tally. Regarding our attack on privacy, the scale and outcome of all known real-world elections based on Helios rule out the possibility of effectively violating the privacy of voters through ballot copying. We also checked the 2010 IACR bulletin board and verified that it does not contain any copied ballot.

Our most realistic new attack challenges the verifiability of elections: we showed that corrupted authorities colluding with a single voter can submit an encryption of an arbitrary (positive or negative) number of approvals for any candidate, and that this encryption is indistinguishable from a normal one. This attack could have a decisive impact on approval elections, where the addition of a reasonable number of votes for a single candidate can easily remain undetected.

Many important Helios elections did not use approval voting, though (e.g., the UCL president election and the IACR 2010 election): in those elections, voters were only allowed to select a limited number of candidates. The capability to submit a single malicious ciphertext has a much more limited effect in that case, due to the need to produce an overall proof of the validity of the ballot besides the individual 0/1 proofs. In this context, two possibilities are left to an attacker: either (1) cheat on an individual proof, that is, if allowed to choose up to n candidates, encrypt n votes for a single candidate, 0 for all others, and the overall proof could still be built normally; or (2) cheat on the overall proof, that

is, select as many candidates as desired and fake the overall proof. The result of these limited manipulations could not have changed the outcome of the two particular elections mentioned above.

Extending our attack to more than a single ciphertext does not seem immediate. Indeed, our attack requires selecting the private key as a function of the hash of all the commitments in one proof. As a result, building two proofs based on different commitments would require using different election keys, which would not be possible in a single election.

Our second most damaging attack relies on authorities that gain access to the randomness that is used by all voters in order to encrypt their messages. This could possibly be achieved by hiding a function that sends this randomness in the JavaScript code sent by the Helios server to the voters for the ballot preparation, or by forcing server-side encryption. Though more demanding, the effect of this attack can also be more severe as the single actively corrupted voter now only needs to submit a regular ballot.

In all cases, including for approval elections (such as the 2011 IACR election) for which our first attack on verifiability applies, there remain possibilities to remove the concerns that our attacks may raise. For instance, a (possibly independent) set of trustees could be asked to run a mixnet on the ciphertexts posted on the bulletin board of the considered elections, which could then be followed by the individual decryption of all shuffled ballots. An invalid ballot would then be detected immediately, and the trustees would not be able to cheat on the decryption of a second ciphertext.

The existence of such a possibility shows that we still are in a better situation than the one obtained with postal voting. Here, the trustees still have a possibility to demonstrate that they did not manipulate the election. That would be much harder for postal voting, where there is no practical way for the tallying officers to demonstrate that the tally they announce actually corresponds to the authentic ballots.

Acknowledgement. We would like to thank Markulf Kohlweiss for helpful comments. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) under the ICT-2 007-216676 European Network of Excellence in Cryptology II and under the HOME/2010/ISEC/AG/INT-011 B-CCENTRE project. Olivier Pereira is a Research Associate of the F.R.S.-FNRS.

References

1. Fiat, A., Shamir, A.: How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987)
2. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Proceedings of the Annual Conference on Computer and Communications Security (CCS). ACM Press (1993)

3. Fouque, P.-A., Pointcheval, D.: Threshold Cryptosystems Secure against Chosen-Ciphertext Attacks. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 351–368. Springer, Heidelberg (2001)
4. Groth, J.: Evaluating Security of Voting Schemes in the Universal Composability Framework. In: Jakobsson, M., Yung, M., Zhou, J. (eds.) ACNS 2004. LNCS, vol. 3089, pp. 46–60. Springer, Heidelberg (2004)
5. Adida, B.: Helios: Web-based open-audit voting. In: Proceedings of the 17th USENIX Security Symposium (Security 2008), pp. 335–348 (2008)
6. Adida, B., de Marneffe, O., Pereira, O., Quisquater, J.J.: Electing a university president using open-audit voting: Analysis of real-world use of helios. In: Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (2009)
7. Helios Headquarters, Princeton University Undergraduate Student Government (2010), <http://usg.princeton.edu/officers/elections-center/helios-headquarters.html>
8. International Association for Cryptologic Research, <http://www.iacr.org/elections/2010>
9. Helios Specification, <http://documentation.heliosvoting.org/verification-specs>
10. Schnorr, C.: Efficient signature generation by smart cards. *Journal of Cryptology* 4, 161–174 (1991)
11. Chaum, D., Pedersen, T.P.: Wallet Databases with Observers. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 89–105. Springer, Heidelberg (1993)
12. Groth, J.: Simulation-Sound NIZK Proofs for a Practical Language and Constant Size Group Signatures. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 444–459. Springer, Heidelberg (2006)
13. Shoup, V., Gennaro, R.: Securing threshold cryptosystems against chosen ciphertext attack. *Journal of Cryptology* 15(2), 75–96 (2002)
14. Tsiounis, Y., Yung, M.: On the Security of ElGamal Based Encryption. In: Imai, H., Zheng, Y. (eds.) PKC 1998. LNCS, vol. 1431, pp. 117–134. Springer, Heidelberg (1998)
15. Schnorr, C.-P., Jakobsson, M.: Security of Signed ElGamal Encryption. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 73–89. Springer, Heidelberg (2000)
16. Wikström, D.: Simplified Submission of Inputs to Protocols. In: Ostrovsky, R., De Prisco, R., Visconti, I. (eds.) SCN 2008. LNCS, vol. 5229, pp. 293–308. Springer, Heidelberg (2008)
17. Bernhard, D., Cortier, V., Pereira, O., Smyth, B., Warinschi, B.: Adapting Helios for Provable Ballot Privacy. In: Atluri, V., Diaz, C. (eds.) ESORICS 2011. LNCS, vol. 6879, pp. 335–354. Springer, Heidelberg (2011)
18. Bulens, P., Giry, D., Pereira, O.: Running mixnet-based elections with helios. In: Shacham, H., Teague, V. (eds.) Electronic Voting Technology Workshop/Workshop on Trustworthy Elections, Usenix (2011)
19. Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations among Notions of Security for Public-Key Encryption Schemes. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 26–45. Springer, Heidelberg (1998)
20. Bellare, M., Rogaway, P.: Entity Authentication and Key Distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994)
21. Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. *Journal of Cryptology* 13(3), 361–396 (2000)

22. Pedersen, T.P.: A Threshold Cryptosystem without a Trusted Party. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 522–526. Springer, Heidelberg (1991)
23. Cortier, V., Smyth, B.: Attacking and fixing helios: An analysis of ballot secrecy. In: CSF, pp. 297–311. IEEE Computer Society (2011)
24. Benaloh, J.: Verifiable Secret-Ballot Elections. PhD thesis, Yale University (January 1987)
25. Pfitzmann, B., Pfitzmann, A.: How to Break the Direct RSA-Implementation of MIXes. In: Quisquater, J.J., Vandewalle, J. (eds.) EUROCRYPT 1989. LNCS, vol. 434, pp. 373–381. Springer, Heidelberg (1990)
26. Naor, M., Yung, M.: Public key cryptosystem secure against chosen ciphertext attacks. In: Proceedings of the Annual Symposium on the Theory of Computing (STOC) 1990, pp. 33–43. ACM Press (1990)
27. Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS) 1999. IEEE Computer Society Press (1999)
28. De Santis, A., Persiano, G.: Zero-knowledge proofs of knowledge without interaction. In: Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS) 1992, pp. 427–436. IEEE Computer Society Press (1992)
29. Cramer, R., Franklin, M.K., Schoenmakers, B., Yung, M.: Multi-authority Secret-Ballot Elections with Linear Work. In: Maurer, U.M. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, pp. 72–83. Springer, Heidelberg (1996)
30. Cramer, R., Gennaro, R., Schoenmakers, B.: A Secure and Optimally Efficient Multi-authority Election Scheme. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 103–118. Springer, Heidelberg (1997)
31. Damgård, I., Groth, J., Salomonsen, G.: The theory and implementation of an electronic voting system. In: Gritzalis, D. (ed.) Secure Electronic Voting. Advances in Information Security, vol. 7, pp. 77–98. Springer (2003)
32. Bernhard, D., Pereira, O., Warinschi, B.: On necessary and sufficient conditions for private ballot submission. IACR Cryptology ePrint Archive 2012, 236 (2012)
33. Smyth, B.: Replay attacks that violate ballot secrecy in helios. IACR Cryptology ePrint Archive 2012, 185 (2012)
34. Küsters, R., Truderung, T., Vogt, A.: Clash Attacks on the Verifiability of E-Voting Systems. In: IEEE Symposium on Security and Privacy (S&P 2012), pp. 395–409. IEEE Computer Society (2012)
35. Camenisch, J.L., Stadler, M.A.: Efficient Group Signature Schemes for Large Groups. In: Kaliski, B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 410–424. Springer, Heidelberg (1997)
36. Chase, M., Kohlweiss, M., Lysyanskaya, A., Meiklejohn, S.: Malleable Proof Systems and Applications. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 281–300. Springer, Heidelberg (2012)