

Adapting Helios for Provable Ballot Privacy

David Bernhard¹, Véronique Cortier², Olivier Pereira³,
Ben Smyth², and Bogdan Warinschi¹

¹ University of Bristol, England

² LORIA - CNRS, France

³ Université Catholique de Louvain, Belgium

Abstract. Recent results show that the current implementation of Helios, a practical e-voting protocol, does not ensure independence of the cast votes, and demonstrate the impact of this lack of independence on vote privacy. Some simple fixes seem to be available and security of the revised scheme has been studied with respect to symbolic models.

In this paper we study the security of Helios using computational models. Our first contribution is a model for the property known as ballot privacy that generalizes and extends several existing ones.

Using this model, we investigate an abstract voting scheme (of which the revised Helios is an instantiation) built from an arbitrary encryption scheme with certain functional properties. We prove, generically, that whenever this encryption scheme falls in the class of *voting-friendly* schemes that we define, the resulting voting scheme provably satisfies ballot privacy.

We explain how our general result yields cryptographic security guarantees for the revised version of Helios (albeit from non-standard assumptions).

Furthermore, we show (by giving two distinct constructions) that it is possible to construct voting-friendly encryption, and therefore voting schemes, using only standard cryptographic tools. We detail an instantiation based on ElGamal encryption and Fiat-Shamir non-interactive zero-knowledge proofs that closely resembles Helios and which provably satisfies ballot privacy.

1 Introduction

Electronic voting protocols have the potential to offer efficient and sound tallying with the added convenience of remote voting. It is therefore not surprising that their use has started to gain ground in practice: USA, Norway and Estonia are examples of countries where e-voting protocols have been, at the very least, trialed in elections on a national scale.

Due to the sensitive nature of elections, security of e-voting protocols is crucial and has been investigated extensively. Among the security properties that have been identified for e-voting, perhaps the most desirable one is that users' votes should remain confidential. Three levels of confidentiality have been identified. These are (in increasing strength) the following.

- Ballot privacy: A voter’s vote is not revealed to anyone.
- Receipt–freeness: A voter cannot obtain information which can prove to a coercer how she voted.
- Coercion resistance: Even a voter who collaborates with a coercer cannot obtain information that proves how she voted.

Other important properties that are desirable include ballot independence [12] (the ballots cast do not depend on each other) and end-to-end verifiability [23,28,38] (it is possible to verify that the election process has been followed honestly).

This paper is motivated by recent developments regarding the security of the Helios voting scheme [45]. Starting from version 2.0 [35], Helios has been using a variant of a classical protocol by Cramer et al. [14] incorporating tweaks proposed by Benaloh [29], and has been used in real-world elections, for example by the International Association for Cryptographic Research (IACR) to elect its 2010 board [36], by Princeton University to elect the undergraduate student government [46] and to elect the president of the Université Catholique de Louvain [35]. Helios aims to achieve only ballot privacy and explicitly discards the stronger confidentiality notions (which it does not satisfy) in favor of efficiency. It turns out that the current implementation of Helios does *not* enforce ballot independence (contrary to the original protocol of Cramer et al. [14]) and, as a result, Cortier and Smyth [37,42] have exhibited several attacks against the ballot privacy property of Helios. (The property is called “ballot secrecy” in Cortier and Smyth’s papers.) The attacks range from simple ballot copying to subtle reuse of parts of existing ballots, however they can all be detected (and prevented) by public algorithms. A revised scheme has been proved secure in a symbolic model but its security in the stronger, computational sense has not been assessed.

Contributions. We start by providing a *computational* security model for ballot privacy (Section 2). In a sense, our model generalizes and strengthens the model of [24,26] where an attacker tries to distinguish when two ballots are swapped. Here, we ask that the adversary cannot detect whether the ballots cast are ballots for votes that the adversary has chosen or not. In doing so, the adversary is allowed to control arbitrarily many players and see the result of the election. Our model uses cryptographic games and thus avoids imposing the more onerous constraints that other definitional styles (in particular simulability) require from protocols.

Next we turn our attention to the revised version of Helios. Our analysis follows a somewhat indirect route: instead of directly analysing the scheme as it has been implemented, we analyze an abstract version of Helios that follows the same basic architecture, but where the concrete primitives are replaced with more abstract versions. Of course, the version we analyze implements the suggested weeding of ballots. We present this abstract scheme as a generic construction of a voting scheme starting from encryption scheme with specific functional and security properties (Section 3).

Focusing on this more abstract version brings important benefits. Firstly, we pin-down more clearly the requirements that the underlying primitives should satisfy. Specifically, we identify a class of *voting-friendly* encryption schemes which when plugged in our construction yield voting schemes with provable ballot privacy. Roughly speaking, such encryption schemes are IND-CCA2 secure and have what we call a homomorphic embedding (parts of the ciphertexts can be seen as ciphertexts of a homomorphic encryption scheme). Secondly, our analysis applies to all voting schemes obtained as instantiations of our generic construction. Although we analyze and propose constructions which for efficiency reasons resort to random oracles, our generic approach also invites other (non-random oracle based) instantiations.

Next, we show how to construct voting-friendly encryption schemes using standard cryptographic tools (Section 4). We discuss two distinct designs. The first construction starts from an arbitrary (IND-CPA) homomorphic encryption scheme and attaches to its ciphertexts a zero-knowledge proof of knowledge of the plaintext. We refer to this construction as the Enc+PoK construction. Despite its intuitive appeal, we currently do not know how to prove that the above design leads to an IND-CCA2 secure encryption scheme (a property demanded by voting-friendliness). We therefore cannot conclude the security of our generic scheme when implemented with an arbitrary Enc+PoK scheme. Nevertheless, an investigation into this construction is important since the instantiation where Enc is the ElGamal scheme and PoK is obtained using the Fiat-Shamir paradigm applied to a Schnorr-like protocol corresponds precisely to the encryption scheme currently used in Helios. The security of this specific construction has been analyzed in prior work. Tsiounis and Yung [17] and Schnorr and Jakobsson [19] demonstrate that the scheme is IND-CCA2 secure, but their proofs rely on highly non-standard assumptions. Nevertheless, in conjunction with the security of our main construction, one can conclude that the current implementation of Helios satisfies ballot privacy based on either the assumption in [17] or those of [19].

We then take a closer look at the Enc+PoK construction and revisit a technical reason that prevents an IND-CCA2 security proof, first studied by Shoup and Gennaro [16]. Very roughly, the problem is that the knowledge extractor associated to the proof of knowledge may fail if used multiple times since its associated security guarantees are only for constant (or logarithmically many) uses. With this in mind, we note that a security proof is possible if the proof of knowledge has a so called *straight line* extractor [22]. This type of extractor can be reused polynomially many times. In this case, the Enc+PoK construction leads to a voting-friendly encryption scheme, whenever Enc is an arbitrary IND-CPA homomorphic encryption scheme.

The second design uses the well-known Naor-Yung transformation [7]. We show that if the starting scheme is an arbitrary (IND-CPA) homomorphic encryption scheme then the result of applying the NY transform is a voting-friendly encryption scheme. Applied generically, the transform may lead to non-efficient schemes (one of its components is a simulation-sound zero-knowledge proof of membership [18]). We present a related construction (where the proof of

membership is replaced by a proof of knowledge) which can be efficiently instantiated in the random oracle model. In the final section of the paper (Section 5) we propose adopting an instantiation of Helios where the encryption-friendly scheme is implemented as above. The computational overhead for this scheme is reasonable (and can be further improved through specific optimization) and the scheme comes with the formal guarantees offered by the results of this paper.

Related work. Chevallier-Mames *et al.* [27] present an unconditional definition of ballot privacy but Helios cannot be expected to satisfy this definition due to its reliance on computational assumptions. Chevallier-Mames additionally show that their definition of unconditional ballot privacy is incompatible with universal verifiability; however, ballot privacy and universal verifiability have been shown to coexist under weaker assumptions, for example as witnessed by Juels, Catalano & Jakobsson [23]. Computational definitions of ballot privacy have been considered by Benaloh *et al.* [2,4,5]. These definitions however do not come with a general characterization of the properties that an encryption scheme should satisfy in order to ensure that they are satisfied (the corresponding security notions did not exist at that time either). Wikström [34] considered the general problem of secure submission of inputs with applications to mixnet-based voting protocols. His definitions and constructions are the most closely related to ours, and will be discussed below. Other definitions for voting systems have been proposed in terms of UC realization of ideal voting functionalities, starting with Groth [21], which capture privacy as part of the functionality behavior.

In addition, receipt-freeness has been considered by Benaloh & Tuinstra [11] and Moran & Naor [25] and coercion resistance has been studied by Juels, Catalano & Jakobsson [23], Küsters, Truderung & Vogt [40] and Unruh & Müller-Quade [39]. These definitions can be used to show ballot privacy because it is believed to be a weaker condition [11,26]; however, they are too strong for protocols which only provide ballot privacy and in particular, they cannot be used to analyse ballot privacy in Helios. Ballot privacy has also been formalized in the symbolic model (for example, [26,33]) but the symbolic model suffers a serious weakness: In general, a correct security proof does not imply the security of the protocol. Cortier & Smyth [37,42] present an attack against ballot privacy in Helios and propose a variant of Helios which aims to prevent the attack by weeding ballots. Their solution has been shown to satisfy ballot privacy in the symbolic model but Cortier & Smyth acknowledge that a thorough cryptographic analysis of the solution is necessary.

2 Ballot Privacy

Notation. Throughout this paper, we use the following notation. Assignment and input/output of algorithms are both denoted by a left-facing arrow \leftarrow . Picking a value x uniformly at random from a set S is denoted by $x \xleftarrow{R} S$. The expression $C \stackrel{\pm}{\leftarrow} c$ appends c to the list C , $()$ on its own is an empty list. We use “C” style returns in algorithms, i.e. “Return $a = b$ ” to mean return 1 if $a = b$, otherwise 0.

A function f is called *negligible* if for any polynomial P , there exists η_0 such that $\forall \eta \geq \eta_0, f(\eta) \leq \frac{1}{P(\eta)}$.

2.1 Voting Schemes

In this section we fix a general syntax for the class of voting schemes that we treat in this paper. In particular, our syntax encompasses several variations of the Helios protocol.

We consider schemes for votes in a non-empty set \mathbb{V} , and we assume \perp to be a special symbol not in \mathbb{V} that indicates that the voter has abstained. The result of an election is then an arbitrary function ρ that takes a list of votes as input and returns the election result. Elections are stateful, so the algorithms that we define next use such a state. Since often, and in particular in the case of Helios, this state is a bulletin board, in the definition below we write BB for this state (and even refer to it as a bulletin board).

Definition 1 (Voting scheme). *Algorithms (Setup, Vote, ProcessBallot, Tally) define a voting scheme as follows.*

- **Setup:** *The setup algorithm takes a security parameter 1^λ as input and returns secret information x , public information y , and initializes the state BB . We write $(x, y, BB) \leftarrow \text{Setup}(1^\lambda)$ for this process. We assume the public information is available to all subsequent algorithms.*
- **Vote:** *The voting algorithm takes a vote $v \in \mathbb{V}$ as input and produces as output a ballot b (that encodes the vote). We write $b \leftarrow \text{Vote}(v)$ for this process.*
- **ProcessBallot:** *The ballot processing algorithm takes a candidate ballot b and a bulletin board BB , checks the ballot for correctness (e.g. that it is well formed, it is not a duplicate, etc.) and returns a result (accept/reject) and the new state of the bulletin board. We write $(a, BB) \leftarrow \text{ProcessBallot}(BB, b)$ for this process. Here a is either accept or reject.*
- **Tally:** *The tallying algorithm takes the secret information x and the bulletin board BB and produces the election result.*

For correctness of the scheme, we demand two conditions: 1) ballot tallying corresponds to evaluating the function ρ on the underlying votes; and 2) correctly constructed votes will be accepted by the ballot processing algorithm. Both conditions should hold with overwhelming probability and can be captured by the experiment described in Figure 1. In this experiment, an adversary repeatedly submits votes $v_1, v_2, \dots \in \mathbb{V}$ and each vote is used to construct a ballot which is then processed. The game outputs 1 (the adversary wins) if the `ProcessBallot` algorithm rejects some ballot or the result of the election does not correspond to the votes cast. The voting scheme is correct if the algorithm outputs 1 with at most negligible probability.

2.2 Security Model

Informally, ballot privacy is satisfied if an adversary in control of arbitrarily many voters cannot learn anything about the votes of the remaining, honest

```

Exp $\Pi$ corr( $A$ )
  ( $x, y, BB$ )  $\leftarrow$  Setup
   $V = ()$ 
  repeat
    ( $a, v$ )  $\leftarrow$   $A$ 
     $b \leftarrow$  Vote( $v$ )
    ( $r, BB$ )  $\leftarrow$  ProcessBallot( $BB, b$ )
     $V \stackrel{\pm}{\leftarrow} v$ 
  until  $a = \text{stop}$  or  $r = \text{reject}$ 
  if  $r = \text{"reject"}$  or Tally( $x, BB$ )  $\neq$   $\rho(V)$  then return 1 else return 0

```

Fig. 1. Experiment for defining the correctness of a voting scheme

voters beyond what can be inferred from the election result. The adversary can read the (public) bulletin board and the communication channels between the honest parties and the bulletin board (in other words, we assume them to be authentic but not secret). Ballot privacy requires that the adversary is unable to distinguish between real ballots and fake ballots, where ballots are replaced by ballots for some fixed vote ε chosen by the adversary.

Formally, we consider an adversary that can issue two types of queries, `vote` and `ballot`, to an oracle \mathcal{O} . The oracle maintains two bulletin boards initialized via the setup algorithm: BB is visible to the adversary and BB' always contains ballots for the real votes. A `vote` query causes a ballot for the given vote to be placed on the hidden BB' . In the real world, the same ballot is placed on BB ; in the fake one a ballot for ε is placed on BB instead. A `ballot` query always causes the submitted ballot to be processed on both boards. This process is defined formally in Figure 2. The experiment on the right of Figure 2 is used to define ballot privacy. The selection of β corresponds to the real world ($\beta = 0$) or the fake world ($\beta = 1$). Throughout the experiment the adversary has access to BB , but tallying is done using BB' .

Definition 2 (Ballot Privacy). *We define the advantage of adversary A in defeating ballot privacy for voting scheme Π by:*

$$\mathbf{Adv}_{\Pi}^{BS}(A) = \Pr[\mathbf{Exp}_{\Pi}^{BS}(A) = 1] - \frac{1}{2}$$

and say that Π ensures ballot privacy if for any efficient adversary its advantage is negligible.

We make a few remarks regarding the security model that we propose. Firstly, we use cryptographic games rather than a simulation based definition. The former offer well-accepted levels of security, are more flexible, and allow for more efficient implementations. Second, we model directly the more relaxed notion of vote privacy and not stronger notions like receipt-freeness or coercion resistance [26]. While stronger notions are certainly desirable, they are more

<pre> vote(v) $b' \leftarrow \text{Vote}(v)$ if $\beta = 0$ then $b \leftarrow b'$ else $b \leftarrow \text{Vote}(\varepsilon)$ $(r, BB) \leftarrow \text{ProcessBallot}(b, BB)$ $(r', BB') \leftarrow \text{ProcessBallot}(b', BB')$ return (r, BB, b) ballot(b) $(r, BB) \leftarrow \text{ProcessBallot}(b, BB)$ if $r = \text{accept}$ then $(r', BB') \leftarrow \text{ProcessBallot}(b, BB')$ return (r, BB) </pre>	<pre> $\text{Exp}_H^{BS}(A)$ $(x, y, BB) \leftarrow \text{Setup}(1^\lambda)$ $BB' \leftarrow BB$ $(\varepsilon, st) \leftarrow A(y)$ $\beta \leftarrow \{0, 1\}$ $st \leftarrow A^\mathcal{O}(st)$ result $\leftarrow \text{Tally}(x, BB')$ $\hat{\beta} \leftarrow A(st, \text{result})$ return $\beta = \hat{\beta}$ </pre>
--	--

Fig. 2. The algorithms on the left explain how the oracle processes adversary’s queries. The experiment on the right is used to define ballot privacy.

difficult to achieve leading to rather inefficient protocols. Indeed, Helios deliberately trades these stronger notions for efficiency. Finally, we emphasize that our computational definition does not mirror existing security definitions in more abstract models, e.g. [24]. It turns out that the direct extension of that definition to computational models seems strictly weaker than the definition that we provide. We comment more on this point later in the paper.

3 A Generic Construction of Voting Schemes with Ballot Privacy

In this section we present a generic construction of a voting scheme starting from any encryption scheme with certain properties. We first fix this class of encryption schemes (which we call voting-friendly), then give our construction and prove its security.

3.1 Voting-Friendly Encryption

In a nutshell, a voting-friendly encryption scheme is a “(threshold) checkable provable IND-CCA2 secure public key encryption scheme with key derivation and a homomorphic embedding”. These rather convoluted looking requirements are in fact not too onerous. We explain informally each of the requirements in turn and give formal definitions. For simplicity, the presentation in this section is for the non-threshold case, that is decryption is carried out using a single key by a single party, as opposed to implementing decryption via an interactive process where several parties share the keys.

Non-Interactive Zero Knowledge Proof Systems. Here we recall some basic notions regarding non-interactive zero-knowledge proof systems [6]. Given language L_R defined by NP relation R we write $(w, x) \in R$ if w is the witness that $x \in L_R$.

A proof system for L_R is given by a pair of algorithms (Prover, Verifier) called prover and verifier, respectively. We distinguish between proof systems in the common reference string model (in this situation, an additional algorithm Setup produces a common reference string accessible to both the prover and the verifier) and the random oracle model (where the setup is not required, but all algorithms in the system have access to a random oracle). In a standard execution of the proof system, the prover and the verifier both have an element $x \in L_R$ as input and in addition, the prover has as input a witness w that $x \in L_R$ (i.e. $R(w, x) = 1$). The prover sends a single message π to the verifier who outputs the decision to accept/reject. We call π a proof for the statement $x \in L_R$. Typical requirements for such proof systems are that they should be sound (if the input x is not in L_R then the verifier rejects π with overwhelming probability) and complete (if x is in the language then the verifier accepts π with probability 1). We write $\pi \leftarrow \text{Prover}$ for the process of producing proof π when the statement x and the witness w are clear from the context. A non-interactive proof system is zero-knowledge if there exists a simulator Sim that is able to produce transcripts indistinguishable from those of a normal execution of the protocol. The simulator may use a trapdoor in the common reference string model, or can program the random oracle in the random oracle model. We occasionally write $(\text{Prover}, \text{Verifier}) : R$ to indicate that the proof system is for the language L_R .

We assume the reader is familiar with public key encryption and its associated security notions. We write $(\text{Gen}, \text{Enc}, \text{Dec})$ for the key generation, encryption, and decryption algorithms of a public key encryption scheme.

Homomorphic encryption. We also briefly recall the notion of *homomorphic* encryption. An encryption scheme is homomorphic if the plaintext space is a group and there exists an algorithm Add that takes two ciphertexts for messages m_0 and m_1 and produces a ciphertext for $m_0 \circ m_1$ (where \circ is the group operation on plaintexts).

Embeddable Encryption. A crucial property for the encryption schemes that are the focus of this section is that they have a *homomorphic embedding*. Informally, this property means that it is possible to identify part(s) of the ciphertexts as forming a ciphertext for some other encryption scheme, and this second encryption scheme is homomorphic. The ElGamal+PoK construction sketched in the previous section is an example of an encryption scheme with an homomorphic embedding. Indeed the e component of a ciphertext (e, π) is a ciphertext for an homomorphic encryption scheme (ElGamal). The next definition makes this discussion more precise.

Definition 3 (Homomorphic Embedding). *We say that the homomorphic encryption scheme $\Pi = (\text{EGen}, \text{EEnc}, \text{EDec}, \text{EAdd})$ is embedded in encryption scheme $\Pi' = (\text{Gen}, \text{Enc}, \text{Dec})$, or alternatively that encryption scheme Π' has Π as a homomorphic embedding if there are algorithms ExtractKey, Extract such that for all m, pk, sk, c*

$$\text{EGen}() = \text{ExtractKey}(\text{Gen}())$$

$$\begin{aligned} \text{EEnc}(m, \text{ExtractKey}(pk)) &= \text{Extract}(\text{Enc}(m, pk)) \\ \text{Dec}(c, sk) &= \text{EDec}(\text{Extract}(c), sk) \end{aligned}$$

Essentially, the `ExtractKey` algorithm maps keys (or key pairs) for the “larger” scheme to keys for the embedded one, and the `Extract` algorithm extracts the ciphertext for the embedded scheme out of ciphertext for the larger one, while performing validity verifications at the same time.

The `Extract` algorithm must, by definition, produce a ciphertext that decrypts to the same value as the input that it is given; in particular it must produce a “ciphertext” that decrypts to \perp if and only if its input does. However, the `Extract` algorithm does not take any secret keys as input. This implies that anyone can check whether a ciphertext is valid (in the sense that it decrypts to something other than \perp) without knowing the secret key. This property forms the basis for combining homomorphic and IND-CCA2 secure encryption in our construction.

We note that an IND-CCA2 secure cryptosystem with homomorphic embedding is actually very close to a submission secure augmented (SSA) cryptosystem as defined by Wikström [34]. Some important differences appear, though. The most important one is that SSA cryptosystems do not require public verifiability of the ciphertexts: it might be necessary to publish a private key augmentation to be able to perform ciphertext validity checks. While this feature enables efficient solutions that are secure in the standard model, it is however often not desirable in practice: it is quite useful to be able to dismiss invalid votes as soon as they are submitted (and to resolve potential conflicts at that time) rather than needing to wait for some partial key to be revealed. Besides, in order to mitigate this inconvenience, SSA cryptosystems allow multiple independent augmentations, which enables updating an augmentation and revealing the previous one in order to be able to check the validity of previously submitted ciphertexts. Our requirement of immediate public verifiability property makes this feature unnecessary for our purpose.

We also note that in concurrent work, Persiano [44] and Smart [41] define similar embedding concepts.

S2P Key Derivation. This property simply requires that if a key pair is produced by the key generation algorithm of an encryption scheme then it is possible to compute the public key from the secret key. This property will allow us to use proofs of knowledge of the secret key corresponding to the public key.

Definition 4 (S2P Key Derivation). *An encryption scheme has the S2P key derivation property if there is an algorithm `DeriveKey` such that $(x, y) \leftarrow \text{Gen}$ implies $y = \text{DeriveKey}(x)$.*

Provable Encryption. In our generic construction voters need to certify that various encryptions in the ballots that they produce satisfy some desirable properties (e.g. that a ciphertext encrypts 0 or 1, and not something else), and such certification can be done via zero-knowledge proofs of knowledge. Since all of the statements that we are interested in are NP statements, the existence of appropriate proof systems follows from general results [9]. Here, we make more precise

the statements for which we demand the existence of such proof systems and introduce some useful notation for the proof systems associated to the various languages that we define.

In particular, it should be possible to prove knowledge of the secret key corresponding to the public key, knowledge of the plaintext underlying a ciphertext, as well as proving that a certain plaintext has been obtained by decrypting with the key associated to the public key. To avoid complex nomenclature, we call a scheme for which this is possible a scheme with provable encryption.

Definition 5 (Provable Encryption). *An encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ is provable if it has the S2P key derivation property and the following non-interactive zero-knowledge proof systems exist:*

1. $(\text{ProveGen}, \text{VerifyGen}): R_1(x, y) := y \stackrel{?}{=} \text{DeriveKey}(x)$
2. $(\text{ProveEnc}, \text{VerifyEnc}): R_2((m, r), c) := c \stackrel{?}{=} \text{Enc}(m; r)$
3. $(\text{ProveDec}, \text{VerifyDec}): R_3(x, (c, y, d)) := y \stackrel{?}{=} \text{DeriveKey}(x) \wedge d \stackrel{?}{=} \text{Dec}(x, c)$

The above definition is for standard encryption schemes. For the case when the encryption scheme that we need is embedded, we demand in addition the existence of proof systems for the following two properties. The first requires that one can prove a statement that involves plaintexts underlying several ciphertexts, and secondly, one should be able to prove that the keys for the embedded schemes in use have been correctly obtained from the keys of the embedding one. This latter condition is a simple adaptation of provability as defined above.

Definition 6 (Provable Embedding). *An encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ for message space M with embedded scheme $(\text{EGen}, \text{EEnc}, \text{EDec})$ has embedded provability for $M' \subseteq M^N$ (for some $N \in \mathbb{N}$) if the following zero-knowledge proof-systems exist:*

1. $(\text{ProveGen}, \text{VerifyGen}): R_4(x, y) := y \stackrel{?}{=} \text{DeriveKey}(x)$
2. $(\text{ProveEnc}, \text{VerifyEnc}): R_5((m_1, m_2, \dots, m_N, r_1, r_2, \dots, r_N), (c_1, c_2, \dots, c_N)) :=$

$$\bigwedge_{i=1}^N c_i \stackrel{?}{=} \text{Enc}(m_i; r_i) \wedge (m_1, \dots, m_N) \in M'$$

3. $(\text{ProveEDec}, \text{VerifyEDec}): R_6(x, (y, d, c)) :=$

$$y \stackrel{?}{=} \text{DeriveKey}(x) \wedge (x', y') \leftarrow \text{ExtractKey}(x, y) \wedge d \stackrel{?}{=} \text{EDec}(x', c)$$

In the last relation, the second conjunct is not a boolean condition, but simply indicates that the keypair (x', y') is derived from (x, y) using the ExtractKey algorithm.

The following definition states all the properties that we require from an encryption scheme in order to be able to implement our generic voting scheme.

Definition 7 (Voting-Friendly Encryption). *A voting-friendly encryption scheme for vote space \mathbb{V} is a public-key scheme for message space M with $\mathbb{V} \subseteq M^N$ such that it is IND-CCA2 secure and has S2P key derivation, an embedded homomorphic scheme and embedded provability for \mathbb{V} .*

Note that voting-friendly encryption requires security guarantees of both the encryption scheme and the contained proof systems.

3.2 Our Generic Construction

In this section we describe a voting scheme based on an arbitrary voting-friendly encryption scheme. The design idea is similar to that of Helios.

The scheme handles elections with multiple candidates. In an election with three candidates a vote is a triple (a, b, c) such that $a, b, c \in \{0, 1\}$ and $a + b + c = 1$. A ballot is then simply formed by individually encrypting each component of the list with an IND-CCA scheme that has an homomorphic embedding, and proving in zero-knowledge that the individual plaintexts in a ballot satisfy the desired relation. To prevent an adversary from casting a vote somehow related to that of an honest voter, we ensure that each ballot cast does not contain any ciphertexts that are duplicates of ones in the ballots already on the bulletin board. This condition is checked while processing ballots.

More formally, denote the set of ciphertexts contained in a ballot b by $\text{Cipher}(b)$ and the set of all ciphertexts on the bulletin board BB by $\text{Cipher}(BB)$, that is $\text{Cipher}(BB) = \bigcup_{b' \in BB} \text{Cipher}(b')$. When submitting a ballot b , we check that $\text{Cipher}(b) \cap \text{Cipher}(BB) = \emptyset$.

Definition 8 (Abstract Voting Scheme). *Let Π be a voting-friendly encryption scheme. The abstract voting scheme $V(\Pi)$ is the construction consisting of algorithms 1–4.*

In our construction, \mathcal{V} is the set of voters, \mathcal{Z} is a party representing “the public” (elements sent to \mathcal{Z} are published) which also functions as a trusted party for generating the initial setup parameters and \mathcal{T} is the trustee of the election (that receives the decryption keys).

If M is the message space of the voting-friendly encryption scheme we consider the space of votes to be $\mathbb{V} \subseteq M^N$ for some $N \in \mathbb{N}$.

We consider result functions of the form $\rho : \mathbb{V}^* \rightarrow M^*$ where $\mathbb{V}^* := \bigcup_{i \in \mathbb{N}_0} \mathbb{V}^i$ (this allows us to tally an arbitrary number of votes) and each component of the range of ρ can be described by a sum of the form $\rho_k = \sum_{i \in \mathbb{N}} a_{i,k} \cdot v_i$ for constants $a_{i,k} \in \mathbb{N}$. This covers the class of result functions that can be computed homomorphically, including normal and weighted sums of votes but also the special case of revealing all the votes and allows us to exploit the homomorphism in the tallying operation: The same operation can be performed on homomorphic ciphertexts using the EAdd algorithm, for which we write \oplus i.e. $a \oplus b := \text{EAdd}(a, b)$. Furthermore, we can define scalar multiplication \otimes on the ciphertexts i.e. $2 \otimes a := \text{EAdd}(a, a)$.

We also provide a public verification algorithm as Algorithm 5 although we do not define this property formally.

Algorithm 1. Setup(1^λ)

 \bar{Z} :

$params \leftarrow \text{Setup}(1^\lambda)$. These parameters are implicitly available to all further algorithms.

 $BB \leftarrow ()$ T : $(x, y) \leftarrow \text{Gen}(1^\lambda)$ $\pi^{Gen} \leftarrow \text{ProveGen}(x, y)$ $\mathcal{Z} \leftarrow (y, \pi^{Gen})$ \mathcal{Z} :

$\text{VerifyGen}(y, \pi^{Gen}) \stackrel{?}{=} 1$ or abort with failure.

Algorithm 2. Vote((v_1, v_2, \dots, v_N))

 $\forall j \in \{1, 2, \dots, N\}$ $c_j \leftarrow \text{Enc}(y, v_j)$ $\pi_j^b \leftarrow \text{ProveEnc}(y, v_j, c_j)$ $b_j \leftarrow (c_j, \pi_j^b)$ output b

Algorithm 3. ProcessBallot(b, BB)

if $\text{VerifyEnc}(b) = 0$ **then return** ("reject", BB) **end if**

for all $c \in \text{Cipher}(b)$ **do**

if $\text{Extract}(c) = \perp$ **then return** ("reject", BB) **end if**

if $\text{Cipher}(b) \cap \text{Cipher}(BB) \neq \emptyset$ **then return** ("reject", BB) **end if**

end for

 $BB \stackrel{\pm}{\leftarrow} b$

return ("accept", BB)

Algorithm 4. Tally(BB)

for all $c_j \in BB$ ($j \in \mathcal{V}$) **do** $e'_j \leftarrow \text{Extract}(c_j)$ **end for**

for all k **do**

$e''_k \leftarrow \bigoplus_{j \in \mathcal{V}} (a_{j,k} \otimes e'_j)$ {I.e. use EAdd to compute ciphertexts for the results.}

$r_k \leftarrow \text{EDec}(x, e''_k)$

$\pi_k^{Dec} \leftarrow \text{ProveEDec}(x, e''_k, r_k)$

end for

 $\mathcal{Z} \leftarrow (r_k, \pi_k^{Dec})_k$

Algorithm 5. Verification

\mathcal{Z} performs the following, aborting if any of the checks (denoted by $\stackrel{?}{=}$) fail. The ordering on \mathcal{V} is a slight abuse of notation; it represents the order the ballots were received in. If successful, the result of the election is r .

```

VerifyGen( $y, \pi^{Gen}$ )  $\stackrel{?}{=} 1$ 
for all  $j \in \mathcal{V}$  do
  ( $c_j, \pi_j^b$ )  $\leftarrow b_j$ 
  VerifyEnc( $b_j$ )  $\stackrel{?}{=} 1$ 
  ( $c_j \notin (c_{j'})_{j' \in \mathcal{V}, j' < j}$ )  $\stackrel{?}{=} 1$ 
   $e'_j \leftarrow \text{Extract}(c_j)$ 
   $e'_j \stackrel{?}{\neq} \perp$ 
end for
 $e' \leftarrow \text{EAdd}(\rho, (e'_j)_{j \in \mathcal{V}})$ 
VerifyEDec( $r, \pi^{Dec}, e'$ )  $\stackrel{?}{=} 1$ 

```

We only prove ballot privacy of our construction formally; correctness follows from the correctness of the voting-friendly encryption scheme. The following theorem states that ballot privacy relies entirely on the security of the underlying voting-friendly scheme.

Theorem 1. *Let Π be a voting-friendly encryption scheme. Then $V(\Pi)$ has ballot privacy.*

To prove the theorem we proceed in two steps. First, we strip the voting scheme of the unnecessary details that concern verifiability, resulting in a scheme that we call “mini-voting”. We prove that ballot privacy for this latter scheme only relies on the IND-CCA2 security of the encryption scheme employed (which highlights IND-CCA2 security as the crucial property needed from the underlying building block). We then explain how to adapt the proof to show the security of $V(\Pi)$.

The full proof can be found in the full version of this paper.

4 Constructions for Voting-Friendly Schemes

In the previous section we gave a generic construction of a voting scheme with ballot privacy starting from an arbitrary voting-friendly encryption scheme. In this section we show that such schemes can be easily constructed using standard cryptographic tools in both the standard and the random oracle models. We discuss three different possibilities.

Encrypt + PoK. This construction does not lead immediately to a voting-friendly scheme but its security is highly relevant to that of Helios, and the design idea forms the basis of a construction that we discuss later.

Under this paradigm, one attempts to construct an IND-CCA2 scheme starting from an IND-CPA scheme and adding to the ciphertext a non-interactive

proof of knowledge of the underlying plaintext. Intuitively, this ensures that an adversary cannot make use of a decryption oracle (since he must know the underlying plaintext of any ciphertext) hence the security of the scheme only relies on IND-CPA security. Unfortunately, this intuition fails to lend itself to a rigorous proof, and currently the question whether Enc+PoK yields an IND-CCA2 scheme is widely open. A detailed treatment of the problem first appeared in [16].

Yet, the question is important for the security of Helios: the current implementation is essentially an instantiation of our generic construction with an Enc+PoK encryption scheme. More precisely the encryption scheme Enc is ElGamal, and the proof of knowledge is obtained by applying the Fiat-Shamir transform to a Schnorr proof. Per the above discussion, no general results imply that the resulting ElGamal+PoK scheme is IND-CCA2 secure (a requirement for voting-friendliness) and our generic result does not apply. However, if one is prepared to accept less standard assumptions, two existing results come in handy. The security of the particular construction that employs ElGamal encryption and Fiat-Shamir zero-knowledge proofs of knowledge has been investigated by Tsionis & Yung [17] and Schnorr & Jakobsson [19]. Both works support the conjecture that the construction is IND-CCA2 but neither result is fully satisfactory. Tsionis & Yung make a knowledge assumption that essentially sidesteps a crucial part in the security proof, whereas the proof of Schnorr & Jakobsson assumes both generic groups [13] and random oracles [10]. Nevertheless, since using either assumption we can show that ElGamal+PoK construction is a voting-friendly scheme, we conclude that Helios satisfies ballot privacy under the same assumptions. Unfortunately, the security of the construction under standard assumptions is a long-standing open question. This observation motivates the search for alternative constructions of voting-friendly schemes.

Straight-line Extractors. To motivate the construction that we discuss now, it is instructive to explain why a proof that Enc+PoK is IND-CCA2 fails. In such a proof, when reducing the security of the scheme to that of the underlying primitive, a challenger would need to answer the decryption queries of the adversary. Since the underlying encryption scheme is only IND-CPA secure, the only possibility is to use the proof of knowledge to extract the plaintext underlying the queried ciphertexts. Unfortunately here the proof gets stuck. Current definitions and constructions for proofs of knowledge only consider single statements and the knowledge extractor works for polylogarithmically many proofs but it may break down (run in exponential time [19]) for polynomially many. Since the IND-CCA2 adversary is polynomially bounded answering all of its decryption queries may thus not be feasible.

A construction that gets around this problem employs a zero-knowledge proof of knowledge with a *straight-line* extractor. Such extractors do not need to rewind the prover and in this case the Enc+PoK construction yields an IND-CCA2 encryption scheme. This notion of extraction and a variation of the Fiat-Shamir transform that turns a sigma-protocol into a non-interactive proof of knowledge with a straight-line extractor in the random oracle model has recently

been proposed by Fischlin [22]. As above, starting with a homomorphic encryption scheme would yield a voting friendly encryption scheme. Unfortunately the construction in that paper is not sufficiently efficient to yield a practical encryption scheme.

The Naor-Yung Transformation. This transformation starts from any IND-CPA secure encryption scheme. An encryption of message m is simply two distinct encryptions c_1 and c_2 of m under the original scheme, together with a simulation-sound zero-knowledge proof π that c_1 and c_2 encrypt the same message with an extra property that we call unique applicability. Formally, we have the following definition.

Definition 9 (Naor-Yung Transformation). *Let $E = (\text{EGen}, \text{EEnc}, \text{EDec})$ be a public-key encryption system. Let $P = (\text{Prove}, \text{Verify}, \text{Sim})$ be a non-interactive zero-knowledge proof scheme for proving (in Camenisch's notation [15])*

$$\text{PoK}\{(m, r_1, r_2) : c_1 = \text{Enc}(y_1, m; r_1) \wedge c_2 = \text{Enc}(y_2, m; r_2)\}$$

with uniquely applicable proofs. Assume the input to Prove is given in the form $(m, y_1, y_2, r_1, r_2, c_1, c_2)$.

The Naor-Yung transformation [7] $NY(E, P)$ of the encryption system is the public-key cryptosystem defined in Algorithm 6.

Algorithm 6. Naor-Yung Transformation

Gen

$(x_1, y_1) \leftarrow \text{EGen}$
 $(x_2, y_2) \leftarrow \text{EGen}$
return $((x_1, x_2), (y_1, y_2))$

Enc $((y_1, y_2), m; (r_1, r_2))$

$c_1 \leftarrow \text{Enc}(y_1, m; r_1)$
 $c_2 \leftarrow \text{Enc}(y_2, m; r_2)$
 $\pi \leftarrow \text{Prove}(m, y_1, y_2, r_1, r_2, c_1, c_2)$
return (c_1, c_2, π)

Dec (c_1, c_2, π)

if Verify $(c_1, c_2, \pi) = 1$ **then return** EDec (x_1, c_2) **else return** \perp **end if**

Sahai [18] showed that the above transformation yields an IND-CCA2 encryption scheme if the starting scheme is IND-CPA and the proof system is simulation-sound and has uniquely applicable proofs (essentially each proof can only be used to prove one statement).

Theorem 2 (Sahai[18]). *If the zero-knowledge proof system P has uniquely applicable proofs then the Naor-Yung transformation $NY(E, P)$ of an IND-CPA secure scheme E gives IND-CCA2 security.*

It turns out that if the starting encryption scheme is homomorphic, then the resulting construction is a voting-friendly encryption scheme. Indeed, the resulting scheme has a homomorphic embedding (given either the first or the second component of the ciphertext) and it is checkable (the checking algorithm only needs to verify the validity of π). As explained earlier, the required proof-systems for provability of the embedding exist, from general results. One can therefore obtain voting schemes with provable ballot privacy in the standard model starting from any homomorphic encryption scheme that is IND-CPA secure in the standard model.

In general, the above construction may not be very efficient (the simulation-sound zero-knowledge proof and associated required proof-systems may be rather heavy). In the random oracle model one can implement the above idea efficiently by replacing the simulation-sound zero-knowledge proof (of membership) with a zero-knowledge proof of knowledge of the message that underlies the two ciphertexts. Interestingly, one may regard the NY transform as providing the underlying encryption scheme with a straight-line extractor (so our previous results already apply).

The following theorem is a variation of the basic Naor-Yung transform applied to our setting.

Theorem 3. *If E is an IND-CPA secure homomorphic encryption scheme with S2P key derivation and P is a zero-knowledge proof of knowledge system with uniquely applicable proofs, then $NY(E, P)$ is a voting friendly encryption scheme.*

5 Application to the Helios Protocol

We propose an enhanced version of Helios 3.0 which is an instantiation of our generic voting scheme with a voting-friendly encryption scheme obtained from ElGamal encryption [1] via the NY transform [7]. The required proof of knowledge is obtained via the Fiat-Shamir transform [3] applied to generalized Schnorr proofs. In this scheme duplicate ballots would be rejected as defined in the ProcessBallot procedure (Algorithm 3). We can further improve the efficiency by reusing some components as described by [20].

Thanks to Theorems 1 and 3, we deduce that the enhanced version of Helios 3.0 (provably) preserves ballot privacy. The modification of Helios we propose does not change the architecture nor the trust assumption of Helios and can be easily implemented. The computational overhead is reasonable (both the length of the messages and the time of computation would at most double and some optimizations can be foreseen). In exchange, we get the formal guarantee that Helios does preserve ballot privacy, a very crucial property in the context of electronic voting. For concreteness, we prove the details of the construction, as well as a proof of security in the full version of this paper.

We emphasize that our results go beyond proving ballot privacy of a particular e-voting protocol. We have identified IND-CCA2 as a sufficient condition for constructing voting schemes satisfying our notion of ballot privacy and have given an abstract construction of a Helios-type voting scheme from IND-CPA secure

homomorphic threshold encryption and non-interactive zero-knowledge proofs of knowledge. Our construction is independent of any hardness assumptions or security models (in particular, the random oracle model). We have formalized the concept of embeddable encryption and showed how to construct IND-CCA2 secure encryption with homomorphic embedding, despite the known impossibility of homomorphic IND-CCA2 secure encryption.

As further work, we plan to extend the definitions and proofs for threshold encryption scheme in order to have a fully complete proof for Helios. We are confident that our proof techniques will apply in a straightforward way. We also wish to investigate the possibility of defining ballot privacy in a more general way, e.g. allowing the current voting algorithm to be replaced by a protocol. Indeed, it could be the case that casting a vote or tallying the vote require more than one step.

Acknowledgements. We are very grateful to Ben Adida for helpful discussions on how to enhance ballot privacy in Helios.

This work was partially supported by the European Commission through the ICT Programme under Contract ICT- 2007-216676 ECRYPT II, by the Interuniversity Attraction Pole P6/26 BCRYPT, and by the European Research Council under the European Unions Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 258865 (ProSecure project). Olivier Pereira is a Research Associate of the Belgian Funds for Scientific Research (F.R.S.-FNRS).

References

1. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory* 31, 469–472 (1985)
2. Cohen (Benaloh), J., Fischer, M.: A Robust and Verifiable Cryptographically Secure Election Scheme. In: *Proceedings of the 26th Symposium on Foundations of Computer Science*, pp. 372–382 (1985)
3. 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)
4. Cohen (Benaloh), J., Yung, M.: Distributing the Power of a Government to Enhance the Privacy of Voters. In: *Proceedings of the 5th Symposium on Principles of Distributed Computing*, pp. 52–62 (1986)
5. Cohen (Benaloh), J.: Verifiable Secret-Ballot Elections. Yale University Department of Computer Science Technical Report number 561 (1987)
6. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications. In: *20th STOC*, pp. 103–112 (1988)
7. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing (STOC 1990)*, pp. 42–437 (1990)
8. Schnorr, C.-P.: Efficient signature generation for smart cards. *Journal of cryptology* 4, 161–174 (1991)

9. Damgård, I.B.: Non-interactive circuit based proofs and non-interactive perfect zero-knowledge with preprocessing. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 341–355. Springer, Heidelberg (1993)
10. Bellare, M., Rogaway, P.: Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In: Proceedings of the 1st ACM Conference on Computer and Communications Security (CCS 1993), pp. 62–73 (1993)
11. Benaloh, J., Tuinstra, D.: Receipt-Free Secret-Ballot Elections. In: Proceedings of the 26th ACM Symposium on Theory of Computing, pp. 544–553 (1994)
12. Gennaro, R.: Achieving independence efficiently and securely. In: Proceedings of the 14th Principles of Distributed Computing Symposium (PODC 1995), pp. 130–136 (1995)
13. Shoup, V.: Lower Bounds for Discrete Logarithms and Related Problems. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (1997)
14. 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)
15. Camenisch, J., Stadler, M.: Efficient group signature schemes for large groups. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 410–424. Springer, Heidelberg (1997)
16. Shoup, V., Gennaro, R.: Securing Threshold Cryptosystems against Chosen Ciphertext Attack. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 1–16. Springer, Heidelberg (1998)
17. 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)
18. Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS 1999), pp. 543–553 (1999)
19. 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)
20. Bellare, M., Boldyreva, A., Staddon, J.: Multi-recipient encryption schemes: Security notions and randomness re-use. In: Desmedt, Y.G. (ed.) PKC 2003. LNCS, vol. 2567. Springer, Heidelberg (2002), <http://cseweb.ucsd.edu/~mihir/papers/bbs.html>
21. 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)
22. Fischlin, M.: Communication-Efficient Non-interactive Proofs of Knowledge with Online Extractors. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 152–168. Springer, Heidelberg (2005)
23. Juels, A., Catalano, D., Jakobsson, M.: Coercion-Resistant Electronic Elections. In: Proceedings of the 4th Workshop on Privacy in the Electronic Society (WPES 2005), pp. 61–70 (2005)
24. Kremer, S., Ryan, M.D.: 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)

25. 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)
26. Delaune, S., Kremer, S., Ryan, M.D.: Coercion-Resistance and Receipt-Freeness in Electronic Voting. In: 19th Computer Security Foundations Workshop (CSFW 2006), pp. 28–42 (2006)
27. Chevallier-Mames, B., Fouque, P., Pointcheval, D., Stern, J., Traoré, J.: On Some Incompatible Properties of Voting Schemes. In: Proceedings of the Workshop on Trustworthy Elections, WOTE 2006 (2006)
28. Participants of the Dagstuhl Conference on Frontiers of E-Voting. Dagstuhl Accord (2007), <http://www.dagstuhlaccord.org/>
29. Benaloh, J.: Ballot Casting Assurance via Voter-Initiated Poll Station Auditing. In: Proceedings of the Second Usenix/ACCURATE Electronic Voting Technology Workshop (2007)
30. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998)
31. Clarkson, M.R., Chong, S., Myers, A.C.: Civitas: Toward a Secure Voting System. In: Proceedings of the 29th Security and Privacy Symposium (S&P 2008), pp. 354–368 (2008)
32. Adida, B.: Helios: Web-based open-audit voting. In: 17th USENIX Security Symposium, pp. 335–348 (2008), http://www.usenix.org/events/sec08/tech/full_papers/adida/adida.pdf
33. Backes, M., Hritcu, C., Maffei, M.: Automated Verification of Remote Electronic Voting Protocols in the Applied Pi-calculus. In: Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF 2008), pp. 195–209 (2008)
34. 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)
35. 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: Proceedings of the 2009 Conference on Electronic Voting Technology/Workshop on Trustworthy Elections (2009)
36. International association for cryptologic research. Election page at <http://www.iacr.org/elections/2010>
37. Cortier, V., Smyth, B.: Attacking and fixing Helios: An analysis of ballot secrecy Website with description and video at <http://www.bensmyth.com/publications/10-attacking-helios/> (Cryptology ePrint Archive, Report 2010/625)
38. 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)
39. Unruh, D., Müller-Quade, J.: Universally Composable Incoercibility. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 411–428. Springer, Heidelberg (2010)
40. Küsters, R., Truderung, T., Vogt, A.: A Game-Based Definition of Coercion-Resistance and its Applications. In: Proceedings of the 23rd IEEE Computer Security Foundations Symposium (CSF 2010), pp. 122–136 (2010)
41. Loftus, J., May, A., Smart, N.P., Vercauteren, F.: On CCA-Secure Fully Homomorphic Encryption, <http://eprint.iacr.org/2010/560>

42. Cortier, V., Smyth, B.: Attacking and fixing Helios: An analysis of ballot secrecy. To appear in: Proceedings of the 24th Computer Security Foundations Symposium, CSF 2011 (2011)
43. Küsters, R., Truderung, T., Vogt, A.: Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study. To appear at the 32nd Security and Privacy Symposium, S&P 2011 (2011) (preprint)
44. Persiano, G.: About the Existence of Trapdoors in Cryptosystems. Work in Progress, <http://libeccio.dia.unisa.it/Papers/Trapdoor/>
45. Helios voting. Website, <http://heliosvoting.org>
46. Helios Headquarters, Princeton University Undergraduate Student Government, <http://usg.princeton.edu/officers/elections-center/helios-headquarters.html>