

REACTIVE: Rethinking Effective Approaches Concerning Trustees in Verifiable Elections

Josh Benaloh¹, Michael Naehrig¹, and Olivier Pereira^{1,2}

¹ Microsoft Research, Redmond, WA, USA

² UCLouvain, B-1348 Louvain-la-Neuve, Belgium

Abstract. For more than forty years, two principal questions have been asked when designing verifiable election systems: how will the integrity of the results be demonstrated and how will the privacy of votes be preserved? Many approaches have been taken towards answering the first question such as use of mixnets and homomorphic tallying. But, in the case of large-scale elections, the second question has always been answered in the same way: decryption capabilities are divided amongst multiple independent “trustees” so that a collusion is required to compromise privacy.

In practice, however, this approach can be fairly challenging to deploy. Even if multiple human trustees are chosen, they typically use software and often also hardware provided by a single voting system supplier, and they rarely have any real opportunity to confirm its correctness. As a result, we observe that trustees are generally not in a position to exercise the independent judgment necessary to ensure privacy.

This Systematization of Knowledge (SoK) paper looks at several aspects of the trustee experience. It begins by surveying and discussing various cryptographic protocols that have been used for key generation in elections, explores their impact on the role of trustees, and notes that even the theory of proper use of trustees is more challenging than it might seem. This is illustrated by showing that one of the only references defining a full threshold distributed key generation (DKG) for elections defines an insecure protocol. Belenios, a broadly used open-source voting system, claims to rely on that reference for its DKG and security proof. Fortunately, it does not inherit the same vulnerability, and we offer a security proof for the Belenios DKG.

The paper then discusses various practical contexts, in terms of humans, software, and hardware, and their impact on the practical deployment of a trustee-based privacy model.

Keywords: Verifiable elections, trustees, distributed key generation, secure hardware.

1 Introduction

The academic approach to preserving privacy in large scale verifiable election protocols is to rely on independent *trustees* who are each responsible for production, maintenance, use, and ultimately destruction of their own cryptographic keys. Examples, in chronological order, include [8,17,2,12,5,11,16,37,32].

Decades of literature, beginning with Benaloh and Yung [8], offer creative cryptographic protocols with excellent properties to achieve precisely this kind of separation during key generation [31,22,15,27]. Such distributed key generation mechanisms are key to enable threshold encryption [19], a well-established technique that provides robustness by allowing a pre-determined number of key-holders to perform the necessary decryptions to complete an election while preventing any set of insufficient size from gaining any information at all.

These approaches are very effective at addressing collusion amongst keyholders and are used both in systems which employ homomorphic tallying and those using modern re-encryption mixnets. Parameters can be set to prevent collusions smaller than a defined size from compromising voter privacy. But this relies upon the assumption that the trustees in an election are technical experts with full agency over their decisions and environment.

In practice, the trustees follow specified protocols for distributed key generation before the election, and must validate that the key they generated is actually used in the election. When encrypted tallies need to be decrypted after the voting period, they must validate a set of ciphertexts to be decrypted and follow protocols for threshold decryption. And, during the whole process, the trustees need to make sure that they remain in control of their secret key material, in order to prevent any abuse. This high level blueprint is common to virtually every verifiable voting system that has been deployed at scale, including Helios [2], Scantegrity II [12], vVote [11], the Verificatum mixnet [37] (used in Estonia for instance), or the Swiss Internet voting system [32].

Human trustees are aided by trustee hardware, computing devices that run software for carrying out the necessary cryptographic operations and securely communicate with the other trustee devices. These protocols are often facilitated by a central authority such as an election administrator that routes the trustees inputs and outputs.

However, ideally, to guarantee independence, trustees should be able to independently evaluate and possibly choose the hardware and software they use: corrupted hardware and software might display everything that a human trustee expects to see while leaking every secret to interested parties. Trusted computing techniques may offer some help but, if they are used to assist trustees, they only help as far as trustees are able to verify that the trusted execution technology is actually in use and is used with the expected code, which may be a challenge of its own.

We observe in this paper that the reality is often quite different. In practice, some or all trustees are often chosen from local communities as representatives of the public, and the key management happens during well orchestrated public ceremonies, using software and hardware provided by authorities. In such a context, trustees do not have the opportunity or the tools to perform any thorough verification steps—even when they happen to have the technical expertise to fulfill such tasks. They are instead required to press the buttons they are told to press on the devices provided to them, which contain pre-loaded software. As a result, there is no real independence and no substantive protection from

curious election administrators who want to view election data that is supposed to remain confidential.

Even when trustees are chosen for their expertise, they are usually required to utilize software and hardware provided to them by an external provider who they may not trust. Furthermore, selecting trustees only from amongst “elites” may restrict their independence and create suspicion from the public.

As a result, the ideal situation of trustees selected as well-trained representatives of the public who utilize hardware and software that they have the means to evaluate as trustworthy seems quite distant from the current practices, and seems to be unlikely to become the rule rather than the exception in the foreseeable future.

Our focus here is on the role of trustees in protecting voter privacy and, in particular, how trustees generate the election keys—similar issues arise around decryption and mixing operations. There are some election systems in which integrity also depends upon trustees (see, for instance, [14,32]). However, the essence of this work is an exploration of how well trustees serve to protect privacy in practice.

While the integrity of the results can potentially be verified over and over by independent parties without access to privileged information and at any chosen time, privacy currently depends on a closed set of trustees being able to complete their tasks diligently. Therefore, without this true independent trustee expertise, and hardware and software independent evaluation, there is very limited basis for confidence in the privacy of the votes, and there is essentially no option to verify that privacy holds.

This paper starts by reviewing protocols for trustees, focusing on distributed key generation (DKG) protocols used in real-world verifiable elections. We discuss the properties of these protocols and the practical demands that they place on the trustees. On our way, we explore technical difficulties in some DKGs and propose a fresh security analysis of the DKG protocol of the Belenios Internet voting system.

We then turn to a discussion of the potential benefits and trade-offs that the reliance on secure hardware can bring. In particular, attestation mechanisms can shift the trust that currently needs to be placed in humans, with limited assessment and verification options, to verifiable attestations produced by hardware that can be independently reviewed and challenged by arbitrary auditors rather than by selected trustees. In such a setting, privacy may become a universally verifiable property rather than a property that depends on a specific set of trustees.

2 Key Generation Protocols

In order to understand the tasks that trustees need to accomplish, we start by describing the most common cryptographic protocols that are used in existing verifiable voting systems.

We focus on distributed key generation (DKG) protocols as a central ingredient of the kind of tasks that trustees are expected to perform. There is of course more in the role of trustees than running a DKG, but the cryptographic operations associated with decryption and mixing bring requirements that are similar to those for key generation, and are described elsewhere, so we do not systematically elaborate on them here. We will discuss post-election operations below.

2.1 Single key

ElGamal encryption [20] is by far the most common algorithm for encrypting votes in verifiable voting systems: it is for instance used in all voting systems mentioned in the introduction above [2,12,11,37,32]. In the simplest form of key generation, the key pair used to encrypt ballots can be generated by one single entity. This approach has been used in Helios 1.0 [1] and, for a long time, in the Estonian Internet voting system, in which all ballots are decrypted using a single key held in an HSM [14]. This solution relies entirely on the security of the HSM device: if the key is somehow extracted from the device, or if someone manages to make that device decrypt non-anonymized encrypted votes, then ballots are not secret anymore. If the HSM device fails and (access to) the decryption key is lost, the election tally cannot be computed. This second concern can be mitigated easily with secret key backups. However, the existence of such backups creates new opportunities for stealing the secret key, increasing the risk of violating vote secrecy.

From a technical point of view, and in the context of ElGamal encryption, the generation process of a single key is extremely simple. A group $\mathbb{G}$ of prime order q is chosen (typically as a public parameter of the election system), together with a generator g of that group. The key generation consists in selecting a secret key $s \leftarrow \mathbb{Z}_q$ uniformly at random, and publishing $K = g^s$ as the encryption public key.

2.2 One-round DKG

To limit the risks of having the single decryption key stolen and potentially used to break the secrecy of all the votes, various voting systems turn to a very simple extension of the above key generation mechanism: n trustees $T_1, \dots, T_n$ independently run the single key generation and publish their public keys. These are then (publicly) combined into a single encryption public key. Here, all n secret keys are needed for decryption.

Concretely, each T_i selects a secret key $s_i \leftarrow \mathbb{Z}_q$ as before and publishes the corresponding $K_i = g^{s_i}$ together with a Schnorr proof of knowledge of s_i in order to prevent so-called rogue key attacks [28]. The encryption public key is computed as $K = \prod_{i=1}^n K_i$, with a corresponding secret key $s = \sum_{i=1}^n s_i$ that is never explicitly computed.

This approach is used in various systems, including Helios [2] (since version 2.0), Belenios [16], and the Swiss Internet voting system [32]. Its use in the context of elections was analyzed by Bernhard et al. [9].

This protocol offers two important benefits. First, it addresses the privacy risk when relying only on a single decryption key: now, in order to break privacy, all n private keys are needed. Second, it keeps most of the simplicity of the single-key protocol. In particular, key generation can still be implemented in a dozen lines of Python code, which can be reviewed easily by a knowledgeable trustee, even during a key generation ceremony.

On the flip side, this protocol is even more fragile than the single-key protocol: losing any one of the n secret keys is sufficient to prevent tallying the election. Again, this can be addressed with backups, which may be less sensitive than in the single-key case since a copy of every secret key is needed to perform decryption operations. Concretely, various solutions have been used: trustees can be paired, and each pair of trustees generates a single secret key of which two copies are kept [10]. In some cases, more elegant solutions can be proposed: for instance, in a setting with 3 trustees sitting around a table, each trustee can give a copy of its secret key to the trustee sitting to its left. This, in effect, offers a two-out-of-three threshold key generation process.

2.3 Threshold DKG

A more general approach to handle the risks of some trustees (or secret keys) being unavailable, is to rely on a threshold distributed key generation protocol, in which only k out of n trustees are needed to perform a decryption operation. As far as we know, the protocols that have been deployed for elections all follow a structure proposed by Pedersen: the n trustees start by generating their key pairs as in the one-round DKG protocol outlined above, then run a verifiable secret sharing (VSS) protocol to share their secret keys s_i with the $n - 1$ other trustees, using a threshold of k [31] (each trustee keeps one share of its own key, so this is k -out-of- n secret sharing). The trustees can combine the shares they received into shares of the secret key s that is the sum of the s_i . The encryption public key K is computed exactly as in the one-round DKG above.

Variations of this protocol are used in various election protocols and have been implemented in software packages, including Verificatum [37], which relies on a protocol proposed by Gennaro et al. [22], Belenios [16], which relies on a protocol by Cortier et al. [15] and ElectionGuard [6], whose DKG protocol is analyzed by Benaloh et al. [7].

The obvious benefit of these protocols is their flexibility for choosing the level of robustness: the original protocol by Pedersen and the one by Gennaro et al. make it possible to choose any $k < n/2$, while the two other protocols aim at supporting any choice of $k \leq n$.

However, they are less convenient. First, they require exchanging encrypted shares between trustees, who are required to all be present at the same time, or to be present multiple times in order to complete the multiple rounds of the protocol. Second, the secret and authentic exchange of shares requires the exchange

and authentication of keys for direct communication between the trustees, which may not be trivial to perform: PKIs are usually not structurally available and may need to be built as part of the DKG. Eventually, the extra protocol complexity also results in extra code complexity: even an expert may be uncomfortable reviewing the code of a full threshold DKG during a key generation ceremony (of course, the review may happen in advance and code hashes can be compared, but this again brings additional complexity).

The additional complexity of a threshold DKG may also impact the election verifier: a verifier for Verificatum and Belenios needs to compute Lagrange coefficients and adjust for the set of trustees present for decryption, which is much less straightforward than in the one-round DKG protocol. ElectionGuard, though, administratively combines all the decryption proofs so that the verification of a decryption operation is as simple as in the single-key case. The focus on simplifying the process of election verification is explicit there.

Eventually, the additional complexity in the key generation also leads to a more complicated security analysis, which we now illustrate in the context of the Belenios DKG.

2.4 The Belenios DKG

Belenios offers two DKG protocols: one that is essentially the single-round protocol above, and a threshold protocol that, on page 2 of the Belenios specification (v. 2.5) [23], is claimed to be “described in [Cortier et al.] [15] and proved in [the works of Cortier et al. and Bernhard et al.] [15,9]”.

As mentioned before, the threshold DKG protocol follows the Pedersen design, which we outline here, given a set of trustees $T_1, \dots, T_n$, a chosen threshold k , and assuming the availability of a public bulletin board that behaves as a broadcast channel.

1. Each trustee T_i chooses a random polynomial $P_i(x) = a_{i,0} + a_{i,1}x + \dots + a_{i,k-1}x^{k-1}$ of degree $k - 1$ with coefficients in $\mathbb{Z}_q$ and posts $K_{i,j} = g^{a_{i,j}}$ for $j \in [0, k - 1]$ on the bulletin board.
2. Each trustee T_i sends $s_{i,j} = P_i(j)$ to every other trustee T_j on a secret and authentic channel.
3. Each trustee T_j verifies the shares it received by checking that $g^{s_{i,j}} = \prod_{\ell=0}^{k-1} (K_{i,\ell})^{(j^\ell)}$ for every value of i . If any check fails, a recovery phase starts, which we do not discuss here.
4. The public key is defined as $K = \prod_{i=1}^n K_{i,0}$ and each trustee T_i computes its decryption key share $z_i = \sum_{j=1}^n s_{j,i}$.

Pedersen proposed this protocol for an honest majority, that is, $n \geq 2k - 1$, and this is the model adopted by Verificatum [37]. However, Cortier et al. offer a proof that, when the protocol is used in combination with ElGamal encryption, it offers IND-CPA security under the DDH assumption for arbitrary threshold choices, that is, for any $k \leq n$ [15]. (The other reference [9] mentioned in the Belenios specification [23] focuses on the one-round non-threshold protocol.)

Pedersen’s DKG is insecure in the case of a dishonest majority. As discussed above, a central aspect of the security of a DKG is to make sure that no subset of less than k trustees can “force” the public key to take a value of their choice, of which they would know the corresponding discrete logarithm. For the single-round protocol (and the ElectionGuard protocol), this is achieved by forcing every trustee to offer a Schnorr proof that it knows the s_i value corresponding to its public key K_i . For the Pedersen protocol, this is achieved by having at least k trustees verify the correctness of the shares they received at Step 3 of the protocol: if k trustees hold correct shares of $a_{i,0}$, then it must be the case that T_i knows $a_{i,0}$. The assumption of an honest majority guarantees that at least k trustees will perform the expected verification steps. However, when there is a dishonest majority, this is not the case anymore (this is overlooked in [15]), and a dishonest trustee might be able to simulate the VSS steps for the honest trustees, while ignoring its secret key share s_i , opening the way for arbitrarily choosing the final public key K .

Similar attacks have been known for a long time (see Langford [28] for instance). Common mitigations include an initial round during which every trustee commits to its $K_{i,j}$ values before opening them and resuming the protocol, or requiring every trustee to provide a Schnorr proof that it knows the discrete logarithms of its $K_{i,j}$ values w.r.t. g . The second option is the one adopted in ElectionGuard [6].

The DKG in Belenios. Fortunately, Belenios does *not* exactly follow Cortier et al. [15] but does something slightly different: it requires each trustee T_j to offer a proof of knowledge of z_j defined as the logarithm in base g of $g^{z_j} = \prod_{i=1}^n g^{s_{i,j}} = \prod_{i=1}^n \prod_{\ell=0}^{k-1} (K_{i,\ell})^{j^\ell}$. To the best of our knowledge, this option is original and has never been publicly analyzed.

We take the opportunity here to offer the first proof of security of the Belenios DKG protocol, which is now the basis of the DKG discussed in version 3.0 of the Belenios specification [24]. We use the name *ElGamal encryption* to designate the traditional ElGamal encryption scheme with single-party key generation, and use *Belenios ElGamal encryption* for the threshold encryption scheme that uses the Pedersen DKG augmented with a Schnorr proof of knowledge of z_j provided by each trustee as its key generation protocol and encrypts with ElGamal encryption. Following other works in the DKG literature [22], we assume here that secret and authentic communication channels for communicating the shares are available between all pairs of trustees, and that trustees verify that they have a common view of the protocol’s public elements.

Theorem 1. *Belenios ElGamal encryption is IND-CPA secure under static corruption of up to $k-1$ trustees ($1 \leq k \leq n$) if ElGamal encryption (with standard single-party key generation) is IND-CPA secure with the same public parameters.*

In other words, any PPT adversary that can break the IND-CPA security of Belenios ElGamal encryption while controlling up to $k-1$ trustees can be efficiently transformed into an IND-CPA adversary against ElGamal encryption.

We observe that this statement makes no assumption regarding the Schnorr protocol that is part of Belenios ElGamal: indeed, this one can be proven to be secure (in the random oracle model) in the presence of a computationally bounded adversary, independently of any computational assumption.

Proof. Let k and n be fixed and define $C = \{1, \dots, k-1\}$ and $H = \{k, \dots, n\}$. Assume, w.l.o.g., that the trustees in the set $\{T_i\}_{i \in C}$ are corrupted, while the others are honest. We design an adversary B against ElGamal encryption that wins the IND-CPA game with a probability equal, up to a negligible difference, to the probability that an adversary A controlling the corrupted trustees breaks the IND-CPA security of Belenios ElGamal encryption.

Given A , we design B as follows: when B receives the ElGamal public parameters $(\mathbb{G}, q, g)$ and the public key K^* from the standard ElGamal challenger, it forwards the public parameters to A . B honestly plays the role of the honest trustees, except for T_n . For emulating T_n , it simulates the sharing of the discrete logarithm of $K_{n,0} = K^*$ by picking $s_{n,i}$ as a random element of $\mathbb{Z}_q$ for every $i \in C$. It then derives $\{K_{n,i}\}_{i \in C}$ so that $g^{s_{n,i}} = \prod_{j=0}^{k-1} (K_{n,j})^{i^j}$ for every $i \in C$, by Lagrange interpolation “in the exponent”—this works because $|C| < k$. The view of A is distributed in a way that is identical to what it would be in a normal execution of the Belenios ElGamal key generation, and A completes the protocol on behalf of the corrupted trustees. When B received all its shares from the corrupted trustees, it can program the random oracle in order to produce simulated Schnorr proofs of knowledge of z_i as the logarithm of $\prod_{j=1}^n g^{s_{j,i}} = \prod_{j=1}^n \prod_{\ell=0}^{k-1} (K_{j,\ell})^{i^\ell}$ in base g for every $i \in H$. B checks the resulting transcripts and fails if anything is wrong.

If all the verification steps succeed, B can now extract the $z_i = \sum_{j=1}^n s_{j,i}$ values for $i \in C$ from the Schnorr proofs provided by A . Extraction from the multiple proofs can be performed because all the proofs are available at the same time. Subtracting the shares that it sent on behalf of the honest trustees, B can also compute $z_{C,i} = \sum_{j \in C} s_{j,i}$ for $i \in C$. When $i \in H$, the values $z_{C,i} = \sum_{j \in C} s_{j,i}$ can also be computed, since the $s_{j,i}$ values are shares that have been sent by A . As a result, B has all n shares of $s_C = \sum_{j \in C} s_{j,0}$, and can reconstruct that value.

Eventually, when A asks for the encryption of a pair of messages (m_0, m_1) , B forwards it to the ElGamal challenger, who returns a ciphertext $(c_0, c_1) = (g^r, m_b(K^*)^r)$. B then submits to A the ciphertext $(c_0, c_1(c_0)^{s_C + \sum_{i=k}^{n-1} s_i}) = (g^r, m_b K^r)$. When A outputs a guess b' on b , B forwards that guess to the ElGamal challenger. The probability that $b = b'$ is exactly the one that A makes a correct guess in the Belenios ElGamal IND-CPA security game. The only possible discrepancy comes from the potential failures to simulate or extract a Schnorr proof, which can be made negligible. $\square$

We summarize the properties of the key generation processes used in the systems we discussed in Table 1.

Table 1. Overview of key generation in various deployed verifiable voting systems.

	Single trustee	n -out-of- n DKG	k -out-of- n DKG (honest maj.)	k -out-of- n DKG (dishonest maj.)
Helios 1.0 [1]	✓			
Estonia [14,38]	✓			
Helios 2.0 [2]		✓		
Belenios [16]		✓		✓
Swiss e-Voting [32]		✓		
Verificatum [37]			✓	
ElectionGuard [6]				✓

3 Protocol setups

The above protocols may be cryptographically correct, but could be useless if their setup assumptions are not satisfied when they are deployed, or if a malicious election administrator gets the possibility to completely circumvent these protocols. For instance, after completing a session of a DKG protocol, trustees need a mechanism to verify that the key used in the election is really the one they generated and has not been replaced with a key that is fully controlled by another entity.

Elections, whether they are electronic or not, cannot exist in isolation: voters need a way to access authentic blank ballots, they need to know where and when to cast their votes, and where to read the election results. This is often achieved by relying on some form of a public bulletin board, and we will assume that such a public bulletin board is available—how to build bulletin boards has been largely discussed in other places [18,26].

The single-key and one-round DKG protocols can directly be implemented when a bulletin board is available: trustees must verify that the public key they produced has been posted on the bulletin board, and that it has not been replaced by the bulletin board manager for instance. Election verifiers can check that all ballots are encrypted and then tallied w.r.t. the correct public keys.

The threshold DKG protocol, however, additionally requires secret and authentic communication channels between trustees to exchange key shares. The assumption that such channels exist is standard and appears virtually everywhere in the DKG literature [31,22].

In many cases, such channels are easy to implement using encryption and signature mechanisms (e.g., via TLS) assuming a trusted public-key infrastructure (PKI). However, in the context of elections, it is much more challenging to depend on a setting in which a trusted PKI exists and certified keys are in the hands of trustees. Trustees will rarely have a certified signature key and, depending on the context, the distributed trust provided by a threshold DKG might be undermined by relying on certificates signed by a single centralized authority, which in turn may be external to the election process and be driven by different incentives.

In practice, this is addressed in various ways. We describe a few examples from systems that use a threshold DKG. These explorations show that this secure point-to-point communication layer, which is just assumed to be available in the DKG literature, is actually far from being obvious to organize in practice.

1. In Verificatum, each trustee generates signing keys and gives the corresponding verification keys to the other trustees. The Verificatum manual suggests on page 2 that, “in practice, the operators could organize a physical meeting to which they bring their laptops and execute the above steps” and “for convenience, hexadecimal encoded hash digests of files can be computed using *vmni* to allow all parties to check that they hold identical protocol info files at the end” [35].
2. The Belenios specification (v.2.5) indicates that each trustee generates signing and encryption keys, which are shared with the other trustees via the voting server. At the opening of the election, each trustee “checks that [its own certificate] appears in the set of verification keys PK of the election” [23].
3. In ElectionGuard, trustees do not generate signing keys. Instead, the ElectionGuard specification (v.2.1) indicates on page 27 that, at the conclusion of the DKG, a preliminary record that contains all public information from the DKG execution is published on the bulletin board. It includes all encryption keys used to exchange shares, the $K_{i,j}$ values, and matching Schnorr proofs. Trustees must hash this information and compare it to a hash computed from their own view of the DKG execution [6].

Verificatum and Belenios create authentic channels using signing keys. Verificatum suggests direct contact between the trustees, allowing direct comparison and confirmation of all the keys that they are going to use. Of course, it remains important for the trustees to verify that the correct key material is also published as part of the official election record on the bulletin board.

In a different approach, Belenios focuses on the direct verification of the election record and requires that trustees confirm the presence of their own signing (and encryption) keys. We pointed out that trustees should agree on all the keys that have been used to ensure that key shares are not compromised—this verification step was absent from Version 2.5 of the specification and, following our comment, a more thorough verification mechanism has been introduced since Version 2.5.1 in order to avoid potential attacks.

ElectionGuard adopts a completely different approach and authenticates the view of the protocol execution rather than using signing keys to sign that view. This has essentially the same effect when verification is successful. Of course, signing keys may additionally provide a non-repudiation property, which could be used for accountability should problems occur. However, when signing key generation is part of the protocol, a malicious trustee might as well complain that the signature verification key published on its behalf is incorrect. Here, the in-person protocol suggested in Verificatum may be advantageous when trustees are required to agree on their keys in person, that is, in a setting where authenticity cannot be questioned. It is also the most demanding option for human trustees.

4 Trustees, their hardware, and their software

To the best of our knowledge, the protocols described in Section 2 guarantee the expected cryptographic security properties. However, as we have seen in Section 3, the associated setup assumption and the context of usage make it so that these security properties cannot be granted by machines *only*: when relying on a designated set of human trustees, these trustees must confirm authenticity of the key material used in the election.

The challenges associated with trustees become particularly pronounced when organizing an election. We elaborate on these issues in this section, again by reviewing different approaches adopted in practice. As before, we focus on the key generation deployment, mainly on the choice of trustees, software, and hardware. The level of detail that we can provide varies a lot, as the available public information is often scarce.

4.1 Key generation ceremonies in practice

Estonia. At least in the early versions of their Internet voting system (pre-IVXV), Estonia used an HSM to generate a single key and decrypt all ballots after a threshold number of election officials inserted their “cryptosticks” [14].

We do not know what is offered for verification here, but we assume that the HSM can produce an attestation that it generated the election key and that the key was never released, as well as a list of all the ciphertexts that were decrypted using that key over its life time.

The task of trustees is highly limited here: vote secrecy essentially depends on proper anonymization of the encrypted votes before they are decrypted (which may not be an obvious operation in the absence of the verifiable mixnet that was introduced in a later version of the system), and on verifying that the HSM has not been abused to perform unauthorized decryption or key export operations.

UCLouvain. Since version 2.0, Helios uses the one-round distributed key generation described above. To the best of our knowledge, Helios has not been deployed in government elections, but we have a description of at least one deployment for a university election at UCLouvain [2].

In that election, trustees were selected from various voter groups and worked with the help of external experts. An in-person meeting was organized for the key generation. Trustees were provided with laptops from which hard-disk drives and network interfaces had been removed. The laptops were booted on Linux using live-CDs, and minimal Python code was provided to generate the election keys, under control of the external experts. Public and secret keys were generated and saved on multiple USB sticks (including copies of the secret keys as backup). The laptops were then turned off and the CDs destroyed. Copies of the public keys were uploaded in the system, and trustees were asked to compare the public keys published on their behalf with their own copies.

So, efforts were made to restrict the ways in which secret keys could be exfiltrated from the hardware, but such measures obviously require expert control.

The software used by the trustees apparently was provided by the same source, but it seems that this code was simple enough to be reviewed during the key generation process—again, this requires expert control. Throughout, trustees remained in control of their secret keys. The potential loss of any one of these keys could have prevented the election organizers from tallying the election—hence the existence of backup copies.

We can see that there are tensions between different incentives here. On the one hand, the cryptographic protocol is designed to put the trustees in control of their keys. On the other hand, should a trustee be missing or lose a key, the tallying process might fail. In such circumstances, the blame is likely to be put more strongly on the election organizer and technology supplier than on the citizen who volunteered to help as a trustee. This actually places a strong incentive on the organizer and technology provider to “tweak” key generation in a way that allows them to obtain copies of all the keys, not to break the privacy of the votes, but to recover from a human failure by a trustee that would badly reflect on them. If the experts are chosen by the organizers, this may undermine the trust that can be granted in the key generation process.

Switzerland. The hardware and roles for key management in the Swiss Post voting system used for Swiss government elections is publicly available [33].

Its DKG is a variation of the one-pass protocol above, with some keys generated by Swiss Post and other keys generated by the cantons organizing elections.

Swiss Post holds four keys, generated on four control components, running four different hardened operating systems in order to mitigate the risks of OS-level exploits (Debian, RedHat, Ubuntu, and Windows are listed). The cantons use multiple laptops deployed in a secure offline environment: they receive the election data through USB sticks.

At the Swiss Post level, the operations are under control of multiple expert teams, even though all of them seem to work under the authority of Swiss Post. The security operations and human resources at the canton level are less documented, and we can guess that they vary among cantons.

At least at Swiss Post, the trustee tasks are performed by experts who would be able to verify the software that they are running, but this comes at the cost of limited independence. The generation of other keys at the Canton level may offer an important level of human independence though. However, the level of independence regarding the software and hardware is not clear: if the voting software, OS, and hardware are provided by Swiss Post, then the effective independence may be reduced. Again, we do not know how this is handled at the canton level.

Franklin County, Idaho. The hardware and key management in a deployment of the ElectionGuard SDK during the 2022 General Election in Franklin County, Idaho is documented [29]. This case is interesting in its own right because ElectionGuard uses a threshold DGK protocol, i.e., trustee devices had to communicate with each other during key generation.

Trustee devices were connected to an administrator device on a local network. Trustees were ordinary citizens who operated devices that they were given, running pre-installed software. At the end of key generation, the trustee devices storing their private keys were kept in safes controlled by the election administrator. The devices themselves required the trustee fingerprint to be activated. However, in order to recover from device failures (and despite the use of a threshold DKG), keys were also exported on thumb drives and stored in the same safe.

The real impact the trustees had here seems minimal. Trustees had no control of the software and hardware, and even the keys were kept in the custody of the election administrator. However, one should not expect that the expertise that can be deployed to run a national voting system in Switzerland can also be deployed in much smaller elections such as those in Franklin County: the election records show that 113 ballots were cast and 2 were spoiled in that election. In that specific deployment, no linkage between the individual paper ballots and the voters who cast them was maintained; so encrypted ballots have a much weaker link to the voter identity than in the Internet voting systems described above, in which a voting server controls the identity of the voter and receives its encrypted ballot. So, even though the security of the keys seems to have been lower in this case, the reliance on the keys for secret ballots seems lower as well.

College Park, Maryland ElectionGuard was also used in College Park, Maryland in 2023—one year after Franklin County, Idaho. The set up and process were very similar—with trustees (termed as “guardians”) selected from the public without special expertise by the city Board of Elections.

A public information session was held on September 27, 2023³ and the tally decryption ceremony was held immediately after the close of the election on November 5, 2025.⁴ The tally decryption ceremony shows how little independent agency the trustees had in the process and how ineffective this approach was at protecting voter privacy from a potentially malicious election administrator.

Summing up. It is clear that, in all these cases, keys that safeguard vote privacy are not in the sole control of trustees that are representatives of the general population concerned about the election.

In all cases, the systems rely on authorities to provide the trustees with experts and trustworthy hardware and software. Even though the entire purpose of splitting keys amongst independent entities is to prevent a curious central authority from simply reading votes, in practice, there is little to stop a curious authority from providing software that reveals keys or performs additional decryptions. And, even if the software is correct, it is unlikely that anyone will notice if a curious authority simply asks for additional decryptions.

³ Available at https://college-park.granicus.com/MediaPlayer.php?view_id=2&clip_id=1405.

⁴ Available at https://college-park.granicus.com/MediaPlayer.php?view_id=2&clip_id=1424 (see, especially, the five-minutes commencing from the ten-minute mark).

Trustees charged with generating and managing keys should be well-trained and thoroughly understand their role. They should bring software and devices obtained from sources they choose to trust. They should inspect the data they are asked to decrypt and ensure that it contains only the values necessary and appropriate to complete the election process. While such assumptions are commonplace and may seem reasonable in theory, practical deployments show that the reality makes this challenging.

5 The Hardware Alternative

The publication of audit data can guarantee the integrity of election results without any requirement to trust designated parties (trustees, election administration, etc.). However, in the solutions described above, even if deployed perfectly, privacy requires trusting that a subset of a fixed group of trustees are behaving correctly. This behavior includes correct human behavior (humans do not distribute copies or misuse their keys) but also correct behavior of software and hardware (the software and the hardware do not leak or misuse secret key material).

Hardware-based security technologies have existed for a long time and have advanced tremendously during the last few years. The capabilities of hardware security modules (HSMs) dramatically expanded, and trusted execution environments (TEEs, including Intel SGX and AMD SEV-SNP) are now readily available. These technologies offer the opportunity to replace vague and difficult to verify assumptions on human processes with clear, specific assumptions on hardware.

HSMs and TEEs can generate keys, log when and how they are used, delete keys when they are not needed anymore, and publish signed attestations of all of these steps. Independent observers can verify the attestations made by these devices and the certificate chain to the device manufacturers. Of course, it is possible that a device does not perform as advertised [4,34] or that the certificate chain to a manufacturer is compromised. But similar assumptions seem to be necessary when trustees need to use authentic software on a secure platform, and these may come on top of human assumptions.

The ability of hardware to offer detailed logging of its use gives external observers evidence that there has been no inappropriate access—something that is not available from human trustees, and that has been already considered for different applications, e.g., offering evidence that a website served specific content [39].

5.1 Secure hardware technologies

Hardware Security Modules. Hardware security modules (HSMs) have been used for decades to provide physical protection for high-value keys that do not need to be used often. They offer tamper-resistance and a variety of means for an authorized set of users to request actions be taken such as generating a key

pair, exporting a public key, and using a protected private key to sign or decrypt given data. While these actions are often limited by default to a specific set of cryptographic algorithms that may not be compatible with the type of DKG protocols described above, custom firmware can often be loaded, and high-end HSMs even offer enclaves that support arbitrary code execution.

These features nicely match the requirements of an election. A key may be generated a month or more before any decryptions are required, and decryptions can be done together in batch in a single session. The principal drawback of HSMs is their operational complexity and high cost, especially if they are required to offer the flexibility and performance needed to execute DKG protocols.

Secure Enclaves. Secure enclaves are a newer technology offered, for example, by Intel’s SGX and AMD’s SEV-SNP technologies. They are capable of executing arbitrary code and providing attestations about the code they ran and the results produced, while offering strong isolation and encrypting all their communication with memory external to the CPU.

Although they are more flexible and far less expensive than HSMs, secure enclaves offer less robust physical protections, as evidenced by an already long history of side-channel vulnerabilities [34]. Another disadvantage of secure enclaves is that they only live as long as their host devices are powered and running. There are some means for preserving and reconstituting state. However, enabling such mechanisms could ruin the security of logging mechanisms that could be used to keep track of every key usage: so-called rollback attacks could be used to perform decryption tasks while escaping logging mechanisms. Generic solutions aiming at detecting rollbacks are being explored [3], but we will discuss more direct options below.

Trusted Platform Modules. Trusted platform modules (TPMs) are security components available in most commercial PCs. Their principal utility is to provide externally-verifiable attestation of the software running on a PC, and they have been successfully deployed to secure open source voting machines [36]. The principal difference between the secure enclave approach and TPMs is that TPMs offer no protection of the data being computed while secure enclaves offer protections to make it more difficult to access data externally. This makes TPM technology less attractive for handling the role of a trustee.

5.2 Resilient Usage of Secure Hardware

What could secure hardware offer? We now outline how secure trusted hardware instances, either permanently instantiated on an HSM, or volatily in a TEE, can be leveraged to offer an alternative to human trustees. The approach that we outline here is radically different from the ones discussed above: there can be a single human trustee, and we will take advantage of the secure hardware instances to make the behavior of the human verifiable by any auditor. So, the trust that is traditionally required to be placed in humans and the software

and hardware they use could now be replaced with trust in hardware only: the rest becomes verifiable by any independent auditor, offering a form of privacy verifiability.

Overview. The starting point is to deploy threshold encryption and decryption protocols on a set of secure hardware instances running on multiple devices. Each instance runs open-source trustee code, and offers an attestation that the expected code is running. The attestations can be tied from the specific hardware components to their manufacturers.

The instances are coordinated by an election administrator. The administrator submits inputs to them and collects their outputs. Crucially, the administrator is required to publish attestations of all operations performed by each instance. A trustee protocol will be defined as secure if the publication of valid attestations produced by the hardware component guarantees the expected security properties, that is, the only decryption operations that will ever be performed w.r.t. the election public key are those attested to by the secure hardware instance.

Intuitively, if the hardware is trusted, a single secure instance would be enough. However, we are concerned that this single device might fail, and the primary reason for using multiple instances is now to offer resilience to hardware failures. This is very different from the primary goal in the context of human trustees, namely to guarantee that a single trustee cannot silently decrypt ballots. Of course, running instances on a heterogeneous set of devices from a variety of vendors also offers additional privacy assurances, should a specific device fail to provide the isolation and attestation properties expected.

Setup phase. To prepare for an election, a setup round generates signature keys that are used to establish secure communication channels between secure instances: each instance generates a signature key pair and submits the signature verification key to the administrator, who returns an election identifier `eid` and the signature verification keys provided by the other instances. Each instance attests to the list of signature verification keys that it will use in the context of election `eid`. This concludes the setup part of the protocol.

The DKG. From this moment on, we can rely on a set of existing protocols as described by Chen and Lindell for instance, who also prove their security [13].

1. *Distributed key generation.* This protocol runs a threshold DKG with a quorum of k within a set of n hardware instances, based on n parallel executions of Feldman’s VSS protocol [21].
2. *Refresh.* This protocol refreshes existing shares of a given secret, essentially by adding to the existing shares a set of freshly generated shares of 0.
3. *Removing a device.* This protocol makes it possible to invalidate the share held by a specific instance, essentially by running a refresh of the shares while excluding that instance from the refresh. At the end of this protocol, the quorum of k is maintained, but only $n - 1$ instances hold shares.

4. *Adding a device.* This protocol makes it possible to add a new device, resuming the level of robustness from a set of $n - 1$ instances to n instances (without any change in the quorum k).

These protocols can be used as follows. The election administrator starts by triggering the execution of the DKG protocol. At the end of this protocol, each hardware instance attests to its public view of the execution and to the resulting public key in particular. Furthermore, each instance also attests to the secure deletion of all the key shares it saw, except for the single one that needs to be kept for future decryption operations.

On a regular basis, the election administrator starts a refresh protocol. This provides proactive security [25] and protects against an adversary who would manage to extract key shares from an increasing number of devices (e.g., by successfully running side-channel attacks). Each refresh resets all the shares and renders the extracted shares useless. This does not need to reflect on decryption proofs, which can be made share independent [6].

At any point in time, an instance might be identified by the administrator as failing—be it rightly so, as a result of a power failure for example, or in order to maliciously exclude the instance. The purpose of such an exclusion might be to elude the requirement for the instance to attest the destruction of its stored key shares, hence weakening the privacy of the votes. However, any such failure statement is required to be followed by an execution of the device removal protocol by the other instances, which will effectively render the share of the excluded device useless.

Of course, such an exclusion reduces the resilience of the system to the failure of other devices: one may progressively reach a state in which more than $n - k$ instances have been removed, preventing any further operation. In order to avoid this situation, a new secure hardware instance must be started, its signature verification key distributed, and the device addition protocol must be executed.

Decryption operations. Whenever tallying operations start, the election administrator submits the required decryption requests to the instances, who simply decrypt whatever they are asked to decrypt without any specific verification—which they could not perform anyway without visibility of the ballots submitted in the election. However, each decryption operation is securely logged by each device.

When the election is complete, the election administrator requires each device to erase its secret key share. Each device eventually provides an attestation of all the decryption operations that have been performed until its final key shares have been erased, and the final erasure is confirmed. All attestations produced by the devices are finally published for auditing.

The core benefit of this approach is that, if the audit succeeds, and under the assumption that a quorum of the devices on which the instances are running offer the advertised security guarantees, then we obtain guarantees on the secrecy of the votes. In effect, we introduce privacy verifiability in elections. Furthermore, should privacy verification steps fail, the election administrator can be identified

as the faulty party, hence offering a form of accountability that may be harder to achieve using traditional approaches, when the origin of an incorrectly decrypted ballot may be harder to identify.

6 The Commitment Alternative

When it is possible, an even more radical alternative is to entirely eliminate trustees by replacing encryption with commitments as in [30]. This can be a viable approach for in-person voting systems in which vote collection devices are used to generate commitments consistent with the selections made by voters.

When in-person ballot collection devices are used to encrypt ballots that will ultimately be published, voter privacy is already dependent upon an assumption that the ballot collection device will not compromise voter privacy. So, it can be very appealing to use only this assumption and eliminate the trustee role entirely. Depending on local jurisdictional rules, each vote collection device can directly publish a proof of its own tally, or tallies and associated proofs can be aggregated by an election administrator who learns nothing more than the sub-tallies of each vote collection device.

This approach has many advantages, but it is more difficult to implement in remote voting scenarios.

7 Conclusions

It has become clear that the realities associated with the use of trustees in elections do not match the benefits derived in theory. We hope that the alternatives described herein will motivate further research that will better align the theory with practice.

Acknowledgments

This paper benefited from many valuable discussions, and we are happy to thank Melissa Chase, Véronique Cortier, Paul England, Pierrick Gaudry, Esha Ghosh, Kim Laine, Jack Richins, Douglas Wikström, and Hervey Wilson for their very helpful comments and inputs.

References

1. Adida, B.: Helios: Web-based open-audit voting. In: Proceedings of the 17th USENIX Security Symposium. pp. 335–348. USENIX Association (2008)
2. Adida, B., de Marneffe, O., Pereira, O., Quisquater, J.: Electing a university president using open-audit voting: Analysis of real-world use of Helios. In: EVT/WOTE '09. USENIX Association (2009)

3. Angel, S., Basu, A., Cui, W., Jaeger, T., Lau, S., Setty, S.T.V., Singanamalla, S.: Nimble: Rollback protection for confidential cloud services. In: 17th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2023. pp. 193–208. USENIX (2023)
4. Bedrune, J.B., Campana, G.: Everybody be cool, this is a robbery! In: IACR Real World Crypto (2020)
5. Benaloh, J., Byrne, M.D., Eakin, B., Kortum, P.T., McBurnett, N., Pereira, O., Stark, P.B., Wallach, D.S., Fisher, G., Montoya, J., Parker, M., Winn, M.: STAR-Vote: A secure, transparent, auditable, and reliable voting system. In: EVT/WOTE '13. USENIX Association (2013)
6. Benaloh, J., Naehrig, M., Pereira, O.: ElectionGuard design specification version 2.1.0. <https://www.electionguard.vote/spec/> (Aug 2024)
7. Benaloh, J., Naehrig, M., Pereira, O., Wallach, D.S.: ElectionGuard: a cryptographic toolkit to enable verifiable elections. In: 33rd USENIX Security Symposium, USENIX Security 2024. USENIX Association (2024)
8. Benaloh, J.C., Yung, M.: Distributing the power of a government to enhance the privacy of voters (extended abstract). In: Proceedings of the Fifth Annual ACM Symposium on Principles of Distributed Computing. pp. 52–62. ACM (1986)
9. Bernhard, D., Pereira, O., Warinschi, B.: How not to prove yourself: Pitfalls of the Fiat-Shamir heuristic and applications to Helios. In: Advances in Cryptology - ASIACRYPT 2012. LNCS, vol. 7658, pp. 626–643. Springer (2012)
10. Bulens, P., Giry, D., Pereira, O.: Running mixnet-based elections with helios. In: 2011 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections, EVT/WOTE '11. USENIX Association (2011)
11. Burton, C., Culnane, C., Schneider, S.A.: vVote: Verifiable electronic voting in practice. *IEEE Secur. Priv.* **14**(4), 64–73 (2016)
12. Carback, R., Chaum, D., Clark, J., Conway, J., Essex, A., Herrnson, P.S., Mayberry, T., Popoveniuc, S., Rivest, R.L., Shen, E., Sherman, A.T., Vora, P.L.: Scantegrity II municipal election at takoma park: The first E2E binding governmental election with ballot privacy. In: 19th USENIX Security Symposium. pp. 291–306. USENIX Association (2010)
13. Chen, Y.H., Lindell, Y.: Feldman’s verifiable secret sharing for a dishonest majority. *Cryptology ePrint Archive*, Paper 2024/031 (2024), <https://eprint.iacr.org/2024/031>
14. Clarke, D., Martens, T.: Real-world Electronic Voting: Design, Analysis and Deployment, chap. E-Voting in Estonia, pp. 129–141. CRC Press (2017)
15. Cortier, V., Galindo, D., Glondou, S., Izabachène, M.: Distributed ElGamal à la Pedersen: Application to Helios. In: Proceedings of the 12th annual ACM Workshop on Privacy in the Electronic Society, WPES 2013. pp. 131–142. ACM (2013)
16. Cortier, V., Gaudry, P., Glondou, S.: Belenios: A simple private and verifiable electronic voting system. In: Foundations of Security, Protocols, and Equational Reasoning - Essays Dedicated to Catherine A. Meadows. LNCS, vol. 11565, pp. 214–238. Springer (2019)
17. Cramer, R., Gennaro, R., Schoenmakers, B.: A secure and optimally efficient multi-authority election scheme. In: Advances in Cryptology - EUROCRYPT 1997. LNCS, vol. 1233, pp. 103–118. Springer (1997)
18. Culnane, C., Schneider, S.A.: A peered bulletin board for robust use in verifiable voting systems. In: IEEE 27th Computer Security Foundations Symposium, CSF 2014. pp. 169–183. IEEE Computer Society (2014)
19. Desmedt, Y., Frankel, Y.: Threshold cryptosystems. In: Advances in Cryptology - CRYPTO '89. LNCS, vol. 435, pp. 307–315. Springer (1989)

20. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Inf. Theory* **31**(4), 469–472 (1985)
21. Feldman, P.: A practical scheme for non-interactive verifiable secret sharing. In: 28th Annual Symposium on Foundations of Computer Science. pp. 427–437. IEEE Computer Society (1987)
22. Gennaro, R., Jarecki, S., Krawczyk, H., Rabin, T.: Secure distributed key generation for discrete-log based cryptosystems. *J. Cryptol.* **20**(1), 51–83 (2007)
23. Glondou, S.: Belenios specification. <https://github.com/glondou/belenios/blob/2.5/doc/specification.tex>, version 2.5
24. Glondou, S.: Belenios specification. <https://github.com/glondou/belenios/blob/3.0/doc/specification.tex>, version 3.0
25. Herzberg, A., Jarecki, S., Krawczyk, H., Yung, M.: Proactive secret sharing or: How to cope with perpetual leakage. In: *Advances in Cryptology - CRYPTO '95*. LNCS, vol. 963, pp. 339–352. Springer (1995)
26. Hirschi, L., Schmid, L., Basin, D.A.: Fixing the achilles heel of e-voting: The bulletin board. In: 34th IEEE Computer Security Foundations Symposium, CSF 2021. pp. 1–17. IEEE (2021)
27. Komlo, C., Goldberg, I.: FROST: flexible round-optimized Schnorr threshold signatures. In: *Selected Areas in Cryptography - SAC 2020*. LNCS, vol. 12804, pp. 34–65. Springer (2020)
28. Langford, S.K.: Weakness in some threshold cryptosystems. In: *Advances in Cryptology - CRYPTO '96*. LNCS, vol. 1109, pp. 74–82. Springer (1996)
29. Microsoft: End-to-end verifiability in real-world elections. <https://www.electionguard.vote/images/EAC%20Report%20Final.pdf> (Jan 2023)
30. Moran, T., Naor, M.: Receipt-free universally-verifiable voting with everlasting privacy. In: *Advances in Cryptology - CRYPTO 2006*. LNCS, vol. 4117, pp. 373–392. Springer (2006)
31. Pedersen, T.P.: A threshold cryptosystem without a trusted party (extended abstract). In: *Advances in Cryptology - EUROCRYPT 1991*. LNCS, vol. 547, pp. 522–526. Springer (1991)
32. Post, S.: Cryptographic primitives of the swiss post voting system. <https://gitlab.com/swisspost-evoting/crypto-primitives/crypto-primitives> (Feb 2024)
33. Post, S.: E-voting architectedocument. https://gitlab.com/swisspost-evoting/e-voting/e-voting-documentation/-/raw/master/System/SwissPost_Voting_System_architecture_document.pdf (February 2024)
34. van Schaik, S., Seto, A., Yurek, T., Batori, A., AlBassam, B., Garman, C., Genkin, D., Miller, A., Ronen, E., Yarom, Y.: SoK: SGX.Fail: How stuff get eXposed. In: *IEEE S&P Symposium* (2024)
35. Verificatum: User manual for the verificatum mix-net. <https://www.verificatum.org/files/vmnum-3.1.0.pdf> (Sep 2022)
36. VotingWorks: Install.md. <https://github.com/votingworks/vxsuite-complete-system/blob/main/INSTALL.md> (Nov 2022)
37. Wikström, D.: Verificatum. <https://www.verificatum.org/> (May 2022)
38. Willemson, J.: Creating a decryption proof verifier for the estonian internet voting system. In: *Proceedings of the 18th International Conference on Availability, Reliability and Security, ARES 2023*. pp. 58:1–58:7. ACM (2023)
39. Zhang, F., Cecchetti, E., Croman, K., Juels, A., Shi, E.: Town crier: An authenticated data feed for smart contracts. In: *Proceedings of the 2016 ACM CCS*. pp. 270–282. ACM (2016)