



ElectionGuard

Design Specification

Version 2.0.0

August 18, 2023

Josh Benaloh and Michael Naehrig

Microsoft Research

Contents

1	Introduction	5
2	Overview	6
3	Components	9
3.1	Parameter requirements	13
3.1.1	Standard baseline cryptographic parameters	13
3.1.2	Parameter base hash	16
3.1.3	Election parameters and the election manifest	16
3.1.4	Election manifest hash	19
3.1.5	Election base hash	19
3.2	Key Generation	20
3.2.1	Overview of key generation	21
3.2.2	Details of key generation	22
3.2.3	Extended base hash	25
3.3	Ballot Encryption	25
3.3.1	Selection encryption	25
3.3.2	Generation of encryption nonces	26
3.3.3	Ballot well-formedness	26
3.3.4	Outline for proofs of ballot correctness	28
3.3.5	Details for proofs of ballot correctness	29
3.3.6	Encryption of contest data	32
3.3.7	Contest data	33
3.3.8	Proof of satisfying the contest selection limit	34
3.4	Confirmation codes	35
3.4.1	Contest hash	35
3.4.2	Confirmation code	35
3.4.3	Ballot chaining	35
3.5	Ballot Aggregation	37
3.6	Verifiable Decryption	37
3.6.1	Partial decryption by available guardians	38
3.6.2	Combination of partial decryptions.	38
3.6.3	Proof of correctness	39
3.6.4	Decryption of contest data	41

3.6.5	Decryption of challenged ballots	43
3.7	The Election Record	45
4	Pre-Encrypted Ballots	47
4.1	Format of Pre-Encrypted Ballots	47
4.1.1	Selection hash	47
4.1.2	Contest hash	48
4.1.3	Confirmation code	48
4.1.4	Short Codes	49
4.2	The Ballot Encrypting Tool	49
4.2.1	Deterministic nonce derivation	50
4.2.2	Using the Ballot Encrypting Tool	50
4.3	The Ballot Recording Tool	51
4.3.1	Using the Recording Tool	51
4.4	The Election Record	52
4.4.1	Election Record Presentation	52
4.5	Verification of Pre-Encrypted Ballots	52
4.6	Hash-Trimming Functions	55
5	Hash Computation	56
5.1	Input data representation	56
5.1.1	Integers modulo the large prime p	56
5.1.2	Integers modulo the small prime q	57
5.1.3	Small integers	58
5.1.4	Strings	58
5.1.5	Files	58
5.2	Hash function	58
5.3	Hashing multiple inputs	59
5.4	Hash function outputs	59
5.5	Domain separation	60
5.5.1	Parameter base, manifest and base hashes	60
5.5.2	Key generation and extended base hash	61
5.5.3	Ballot encryption and confirmation codes	62
5.5.4	Verifiable decryption	63
5.5.5	Pre-encrypted ballots	63

6	Verifier Construction	64
6.1	Implementation Details	64
6.2	Validation Steps	65
6.2.1	Parameter verification	65
6.2.2	Key ceremony verification	67
6.2.3	Extended base hash verification	68
6.2.4	Ballot verification	69
6.2.5	Tally verification	73
6.2.6	Verification of correct contest data decryption	76
6.2.7	Verification of challenged ballots	77
6.2.8	Verification of pre-encrypted ballots	80
7	Applications to End-to-End Verifiability and Risk-Limiting Audits	84

1 Introduction

This document describes the cryptographic details of the design of ElectionGuard which can be used in conjunction with many new and existing voting systems to enable both end-to-end (E2E) verifiability and privacy-enhanced risk-limiting audits (RLAs). ElectionGuard is not a complete election system. It instead provides components that are designed to be flexible and to promote innovation by election officials and system developers. When properly used, it can promote voter confidence by empowering voters to independently verify the accuracy of election results.

End-to-end (E2E) verifiability

An E2E-verifiable election provides artifacts which allow voters to confirm that their votes have been accurately recorded and counted. Specifically, an election is End-to-end (E2E) verifiable if two properties are achieved.

1. Individual voters can verify that their votes have been accurately recorded.
2. Voters and observers can verify that all recorded votes have been accurately counted.

An E2E-verifiable tally can be used as the primary tally in an election or as a verifiable secondary tally alongside traditional methods. ElectionGuard is compatible with in-person voting – either using an electronic ballot-marking device or an optical scanner capable of reading hand-marked or machine-marked ballots, with voting by mail, and even with Internet voting¹.

Risk-limiting audits (RLAs)

RLAs offer election administrators efficient methods to validate reported election tallies against physical ballot records. There are several varieties of RLAs, but the most efficient and practical are *ballot-comparison audits* in which electronic *cast-vote records* (CVRs) are individually compared against physical ballots.

The challenge with ballot-comparison audits is that public release of the full set of CVRs can compromise voter privacy while an audit without public disclosure of CVRs offers no basis for public confidence in the outcome. ElectionGuard can bridge this gap by enabling public disclosure of encrypted ballots that can be matched directly to physical ballots selected for auditing and can also be proven to match the reported tallies.

About this specification

This specification can be used by expert reviewers to evaluate the details of the ElectionGuard process and by independent parties to write ElectionGuard implementations. This document is not intended to be a fully detailed implementation specification. It does not specify serialization and data structures for the election record or mappings between the notation used here and the corresponding data types and components of a specific implementation. However, this document, together with a

¹Note that there are many challenges to responsible Internet voting that are mitigated but not fully solved by E2E-verifiability. The 2015 U.S. Vote Foundation report at <https://www.usvotefoundation.org/E2E-VIV> details many of these issues, and the 2018 National Academies report at <https://nap.nationalacademies.org/catalog/25120/securing-the-vote-protecting-american-democracy> includes a section on Internet voting (pp. 101–105).

detailed implementation specification or a well-documented ElectionGuard implementation, can be used by independent parties to write ElectionGuard verifiers to confirm the consistency of election artifacts with announced election results.

2 Overview

This section gives a very brief and high-level overview of the ElectionGuard system’s functionality and how it can be used during an election or in a post-election audit such as an RLA.

To use ElectionGuard in an election, a set of *guardians* is enlisted to serve as trustees who manage cryptographic keys. The members of a canvassing board can serve as guardians. Prior to the commencement of voting or auditing, the guardians work together to form a public encryption key that will be used to encrypt individual ballots.

After the conclusion of voting or auditing, a *quorum* of guardians is necessary to produce the artifacts required to enable public verification of the tally.

Key generation

Prior to the start of voting (for an E2E-verifiable election) or auditing (for an RLA), the election guardians participate in a process wherein they generate public keys to be used in the election. Each guardian generates its own public-secret key pair. These public keys will be combined to form a single public key with which votes are encrypted. They are also used individually by guardians to exchange information about their secret keys so that the election record can be produced after voting or auditing is complete – even if not all guardians are available at that time.

The key generation ceremony begins with each guardian publishing its public keys together with proofs of knowledge of the associated secret keys. Once all public keys are published, each guardian uses each other guardian’s public key to encrypt shares of its own secret keys. Finally, each guardian decrypts the shares it receives from other guardians and checks them for consistency. If the received shares verify, the receiving guardian announces its completion. If any shares fail to verify, the receiving guardian challenges the sender. In this case, the sender is obliged to reveal the shares it sent. If all challenged guardians release their challenged shares and these shares all verify, the ceremony concludes and the election proceeds. If a challenged guardian fails to produce key shares that verify, that guardian is removed and the key generation ceremony restarts with a replacement guardian.

Ballot encryption

In most uses, the election system makes a single call to the ElectionGuard API after each voter completes the process of making selections or with each ballot to be encrypted for an RLA. ElectionGuard will encrypt the selections made by the voter and return a confirmation code which the system should give to the voter.²

This is the only point where an existing election system must interface with ElectionGuard. In most uses of ElectionGuard, voters will have an opportunity to challenge their encrypted ballots and

²The confirmation code is not necessary for RLA usage.

view their decryptions to ensure that the encryptions are correct.³ In certain vote-by-mail scenarios and when ElectionGuard is used within an RLA, cast-vote records can be provided in batch without any interface between the voting equipment and ElectionGuard.

The encrypted ballots are published along with non-interactive zero-knowledge (NIZK) proofs of their well-formedness. The proofs assert that an encrypted ballot is well-formed, which means that it is a legitimate ballot and adheres to the limits imposed on selection options and contests. For example, they prove that a selection did not receive more votes than allowed and that no more than the allowed number of votes were received across the selection options in each contest. The encryption method used herein has a homomorphic property which allows the encrypted ballots to be combined into a single aggregate ballot which consists of encryptions of the election tallies.

Verifiable decryption

In the final step, election guardians independently use a share of the secret decryption key, which each guardian computes from the previously shared secret key fragments, to jointly decrypt the election tallies and generate associated verification data. It is not necessary for all guardians to be available to complete this step. If some guardians are missing, a quorum of guardians is sufficient to complete decryption and generate the verification data.

Independent verification

Observers can use this open specification and/or accompanying materials to write independent *election verifiers* that can confirm the well-formedness of each encrypted ballot, the correct aggregation of these ballots, and the accurate decryption of election tallies.

The verification of an election consists of various verification steps that can be separated into three groupings.

Verification of key generation. The first group of verification steps pertain to the key generation process. Under normal circumstances, guardians will engage in a key generation ceremony and produce an election public key without controversy, and there is no need for independent verification of this process. However, if one or more parties to the key generation ceremony complain about the process, the associated independent verification steps can help to resolve conflicts.

These steps should therefore be considered to be optional. The integrity of an election record can be fully verified by independent parties without any of the key verification steps. However, a “premium” verifier may wish to include these steps to verify the correctness of guardian actions during the key generation ceremony.

The key generation verification steps are highlighted in orange boxes.

³A ballot that has been decrypted should be regarded as a test ballot and should not be included in an election tally. After a ballot is challenged, the voter should have an opportunity to cast a fresh ballot. In an E2E-verifiable election, a decrypted ballot should never be cast. However, in an RLA, some anonymized cast ballots may ultimately be challenged and decrypted.

Verification of ballot correctness. As will be described in detail below, there are instances in both the E2E-verifiability application and the RLA application where it is desirable to verifiably decrypt individual ballots. In the former application, this allows voters to confirm that their selections have been correctly recorded within an encrypted ballot, and in the latter application, this allows encrypted electronic ballots that have been selected for audit to be compared with their associated paper ballots.

The steps required to verify the correct decryption of a ballot are grouped together to allow independent verifiers to be easily constructed whose sole purpose is to allow voters or observers to directly verify correct decryption of individual ballots.

The ballot correctness verification steps are highlighted in blue boxes.

Verification of election record. The remaining verification steps apply to the election as a whole. These steps enable independent verification that every ballot in an election record is well-formed, that cast ballots have been correctly aggregated homomorphically, and that these aggregated values have been correctly decrypted to produce election tallies.

It is critical that a complete verification of an election record also includes verification of the correct decryption of any individual ballots that have been decrypted as part of the process. A complete election verifier must therefore incorporate an individual ballot generator as described above.

The election record verification steps are highlighted in green boxes.

Using this specification

The principal purposes of this document are to specify the functionality of the ElectionGuard toolkit and to provide details necessary for independent parties to write election verifiers that consume the artifacts produced by the toolkit.

3 Components

This section describes the four principal components of ElectionGuard.

1. *Parameter Requirements*: These are properties required of parameters to securely and efficiently instantiate the cryptographic components of ElectionGuard. These *cryptographic parameters* are fixed ahead of every election and the same parameters can be used across multiple elections. A specific, recommended set of standard parameters is provided. In addition, there are properties required of the *election parameters* that define the election contests, selectable options, and ballot styles, among other information.
2. *Key Generation*: Prior to each individual election, guardians must generate individual public-secret key pairs and exchange shares of secret keys to enable completion of an election even if some guardians become unavailable. Although it is preferred to generate new keys for each election, it is permissible to use the same keys for multiple elections so long as the set of guardians remains the same. A complete new set of keys must be generated if even a single guardian is replaced.
3. *Ballot Encryption*: While encrypting the contents of a ballot is a relatively simple operation, most of the work of ElectionGuard is the process of creating externally-verifiable artifacts to prove that each encrypted ballot is well-formed (i.e., its decryption is a legitimate ballot without overvotes or improper values).
4. *Verifiable Decryption*: At the conclusion of each election, guardians use their secret key shares to jointly produce election tallies together with verifiable artifacts that prove that the tallies are correct.

Notation

In the remainder of this specification, the following notation will be used.

- $\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$ is the set of integers.
- $\mathbb{Z}_n = \{0, 1, 2, \dots, n - 1\}$ is the ring of integers modulo n .
- \mathbb{Z}_n^* is the multiplicative subgroup of \mathbb{Z}_n that consists of all invertible elements modulo n . When p is a prime, $\mathbb{Z}_p^* = \{1, 2, 3, \dots, p - 1\}$.
- \mathbb{Z}_p^r is the set of r -th-residues in \mathbb{Z}_p^* . Formally, $\mathbb{Z}_p^r = \{y \in \mathbb{Z}_p^* \text{ for which there exists } x \in \mathbb{Z}_p^* \text{ such that } y = x^r \text{ mod } p\}$. When p is a prime for which $p - 1 = qr$ with q a prime that is not a divisor of the integer r , then \mathbb{Z}_p^r is an order- q cyclic subgroup of \mathbb{Z}_p^* and for each $y \in \mathbb{Z}_p^*$, $y \in \mathbb{Z}_p^r$ if and only if $y^q \text{ mod } p = 1$.
- $x \equiv_n y$ is the predicate that is true if and only if $x \text{ mod } n = y \text{ mod } n$.
- The function $\text{HMAC}(\cdot, \cdot)$ shall be used to denote the HMAC-SHA-256 keyed Hash Message Authentication Code (as defined in NIST PUB FIPS 198-1⁴) instantiated with SHA-256 (as defined in NIST PUB FIPS 180-4⁵). HMAC takes as input a key k and a message m of arbitrary length and returns a bit string $\text{HMAC}(k, m)$ of length 256 bits.
- The ElectionGuard hash function $H(\cdot, \cdot)$ is instantiated with HMAC and thus has two inputs, both given as byte arrays. The first input to H is used to bind hash values to a specific

⁴NIST (2008) The Keyed-Hash Message Authentication Code (HMAC). In: FIPS 198-1. <https://csrc.nist.gov/publications/detail/fips/198/1/final>

⁵NIST (2015) Secure Hash Standard (SHS). In: FIPS 180-4. <https://csrc.nist.gov/publications/detail/fips/180/4/final>

election. The second is a byte array of arbitrary length and consists of domain separation bytes and the data being hashed. H outputs a digest of 256 bits, which is interpreted as a byte array of length 32. The detailed specification for H is given in Section 5 below.

- The symbol \oplus denotes bitwise XOR.
- The symbol \parallel denotes concatenation.
- In general, the variable pairs (α, β) , (a, b) , and (A, B) will be used to denote encryptions. Specifically, (α, β) will be used to designate encryptions of votes (usually an encryption of a zero or a one), (A, B) will be used to denote aggregations of encryptions – which may be encryptions of larger values, and (a, b) will be used to denote encryption commitments used to prove properties of other encryptions.

Encryption of votes

Encryption of votes in ElectionGuard is performed using a public key encryption method suggested by Devillez, Pereira, and Peters,⁶ which is a variant exponential form of the ElGamal cryptosystem,⁷ which is, in turn, a static form of the widely-used Diffie-Hellman(-Merkle) key exchange⁸. This encryption method is called *DPP vote encryption* or simply *vote encryption* in this document and rests on precisely the same security basis as Diffie-Hellman(-Merkle) key exchange—which is used to protect the confidentiality of the vast majority of Internet traffic.

Primes p and q are publicly fixed such that q is not a divisor of $r = (p - 1)/q$. A generator g of the order q subgroup \mathbb{Z}_p^r is also fixed. (Any $g = x^r \bmod p$ for which $x \in \mathbb{Z}_p^*$ suffices so long as $g \neq 1$.)

A public-secret key pair can be chosen by selecting a random $s \in \mathbb{Z}_q$ as a secret key and publishing $K = g^s \bmod p$ as the corresponding public key⁹.

A message $m \in \mathbb{Z}_p^r$ is then encrypted by selecting a random nonce $\xi \in \mathbb{Z}_q$ and forming the pair $(\alpha, \beta) = (g^\xi \bmod p, K^m \cdot K^\xi \bmod p) = (g^\xi \bmod p, K^{m+\xi} \bmod p)$. An encryption (α, β) can be decrypted by the holder of the secret s as $\beta/\alpha^s \bmod p = K^m \bmod p$ because

$$\frac{\beta}{\alpha^s} \equiv_p \frac{K^m \cdot K^\xi}{(g^\xi)^s} \equiv_p \frac{K^m \cdot (g^s)^\xi}{(g^\xi)^s} \equiv_p \frac{K^m \cdot g^{\xi s}}{g^{\xi s}} \equiv_p K^m. \quad (1)$$

The value of m can be computed from $K^m \bmod p$ as long as the message m is limited to a small, known set of options.¹⁰

⁶Devillez H., Pereira, O., Peters, T. (2022) *How to Verifiably Encrypt Many Bits for an Election?* in ESORICS 2022 Lecture Notes in Computer Science, vol 13555. Springer, Berlin, Heidelberg. https://link.springer.com/chapter/10.1007/978-3-031-17146-8_32

⁷ElGamal T. (1985) *A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms*. In: Blakley G.R., Chaum D. (eds) *Advances in Cryptology*. CRYPTO 1984. Lecture Notes in Computer Science, vol 196. Springer, Berlin, Heidelberg. https://link.springer.com/content/pdf/10.1007/3-540-39568-7_2.pdf

⁸Diffie W., Hellman, M. (1976) *New Directions in Cryptography* IEEE Transactions on Information Theory, vol 22

⁹As will be seen below, the actual public key used to encrypt votes will be a combination of separately generated public keys. So, no entity will ever be in possession of a secret key that can be used to decrypt votes.

¹⁰The simplest way to compute m from $K^m \bmod p$ is an exhaustive search through possible values of m . Alternatively, a table pairing each possible value of $K^m \bmod p$ with m can be pre-computed. A final option which can accommodate a larger space of possible values for m is to use Shanks's baby-step giant-step method as described in the 1971 paper *Class Number, a Theory of Factorization and Genera*, Proceedings of Symposium in Pure Mathematics, Vol. 20, American Mathematical Society, Providence, 1971, pp. 415-440.

The value of K^m and therefore m can also be computed from the encryption nonce ξ , namely via $\beta/K^\xi \equiv_p K^m$. While the secret key s allows decryption of any ciphertext encrypted to the public key K , the encryption nonce only allows decryption of the specific ciphertext it was used to generate. The encryption nonce must therefore be securely protected. Release of an encryption nonce can, when appropriate, serve as a fast and convenient method of verifiable decryption.

Usually, only two possible messages are encrypted in this way by ElectionGuard. An encryption of one is used to indicate that an option is selected, and an encryption of zero is used to indicate that an option is not selected. For some voting methods such as cumulative voting, Borda count, and a variety of cardinal voting methods like score voting and STAR-voting, it can be necessary to encrypt other small, positive integers.

Homomorphic properties

A fundamental quality of the DPP vote encryption is its additively homomorphic property. If two messages m_1 and m_2 are respectively encrypted as $(A_1, B_1) = (g^{\xi_1} \bmod p, K^{m_1+\xi_1} \bmod p)$ and $(A_2, B_2) = (g^{\xi_2} \bmod p, K^{m_2+\xi_2} \bmod p)$, then the component-wise product

$$(A, B) = (A_1 A_2 \bmod p, B_1 B_2 \bmod p) = (g^{\xi_1+\xi_2} \bmod p, K^{(m_1+m_2)+(\xi_1+\xi_2)} \bmod p) \quad (2)$$

is an encryption of the sum $m_1 + m_2$. (There is an implicit assumption here that $(m_1 + m_2) < q$ which is easily satisfied when m_1 and m_2 are both small. If $(\xi_1 + \xi_2) \geq q$, $(\xi_1 + \xi_2) \bmod q$ may be substituted without changing the equation since $g^q \equiv_p 1$.)

This additively homomorphic property is used in two important ways in ElectionGuard. First, all of the encryptions of a single option across ballots can be multiplied to form an encryption of the sum of the individual values. Since the individual values are one (or some other integer for certain voting methods) on ballots that select that option and zero otherwise, the sum is the tally of votes for that option and the product of the individual encryptions is an encryption of the tally.

The other use is to sum all of the selections made in a single contest on a single ballot. In the simplest case, after demonstrating that each option is an encryption of either zero or one, the product of the encryptions indicates the number of options that are encryptions of one, and this can be used to show that no more ones than permitted are among the encrypted options – i.e., that no more options were selected than permitted. When larger integers are allowed, i.e. an option can receive multiple votes or weighted votes, the product of the ciphertexts then encrypts the total number of votes or the sum of weights and is used in the same way to ensure only the permitted number of votes or permitted sum of weights were used.

However, as will be described below, it is possible for a holder of a nonce ξ to prove to a third party that a pair (α, β) is an encryption of m without revealing the nonce ξ and without access to the secret s .

Non-interactive zero-knowledge (NIZK) proofs

ElectionGuard provides numerous proofs about encryption keys, encrypted ballots, and election tallies using the following four techniques.

1. A Schnorr proof¹¹ allows the holder of a secret key s to interactively prove possession of s without revealing s .
2. A Chaum-Pedersen proof¹² allows an encryption to be interactively proven to decrypt to a particular value without revealing the nonce used for encryption or the secret decryption key s . (This proof can be constructed with access to either the nonce used for encryption or the secret decryption key.)
3. The Cramer-Damgård-Schoenmakers technique¹³ enables a disjunction to be interactively proven without revealing which disjunct is true.
4. The Fiat-Shamir heuristic¹⁴ allows interactive proofs to be converted into non-interactive proofs.

Using a combination of the above techniques, it is possible for ElectionGuard to demonstrate that keys are properly chosen, that ballots are well-formed, and that decryptions match claimed values.¹⁵

Threshold encryption

Threshold encryption is used for encryption of ballots and other data. This form of encryption makes it very easy to combine individual guardian public keys into a single public key. It also offers a homomorphic property that allows individual encrypted votes to be combined to form encrypted tallies.

The guardians of an election will each generate a public-secret key pair. The public keys will then be combined (as described in the following section) into a single election public key which is used to encrypt all selections made by voters in the election.

At the conclusion of the election, each guardian will compute a verifiable partial decryption of each tally. These partial decryptions will then be combined to form full verifiable decryptions of the election tallies.

To accommodate the possibility that one or more of the guardians will not be available at the conclusion of the election to form their partial decryptions, the guardians will cryptographically share¹⁶ their secret keys amongst each other during key generation in a manner to be detailed in Section 3.2. Each guardian can then compute a share of the secret decryption key, which they

¹¹Schnorr C.P. (1990) Efficient Identification and Signatures for Smart Cards. In: Brassard G. (eds) *Advances in Cryptology — CRYPTO' 89 Proceedings*. CRYPTO 1989. Lecture Notes in Computer Science, vol 435. Springer, New York, NY. https://link.springer.com/content/pdf/10.1007%2F0-387-34805-0_22.pdf

¹²Chaum D., Pedersen T.P. (1993) *Wallet Databases with Observers*. In: Brickell E.F. (eds) *Advances in Cryptology — CRYPTO' 92*. CRYPTO 1992. Lecture Notes in Computer Science, vol 740. Springer, Berlin, Heidelberg. https://link.springer.com/content/pdf/10.1007%2F3-540-48071-4_7.pdf

¹³Cramer R., Damgård I., Schoenmakers B. (1994) *Proofs of Partial Knowledge and Simplified Design of Witness Hiding Protocols*. In: Desmedt Y.G. (eds) *Advances in Cryptology — CRYPTO' 94*. CRYPTO 1994. Lecture Notes in Computer Science, vol 839. Springer, Berlin, Heidelberg. https://link.springer.com/content/pdf/10.1007%2F3-540-48658-5_19.pdf

¹⁴Fiat A., Shamir A. (1987) *How To Prove Yourself: Practical Solutions to Identification and Signature Problems*. In: Odlyzko A.M. (eds) *Advances in Cryptology — CRYPTO' 86*. CRYPTO 1986. Lecture Notes in Computer Science, vol 263. Springer, Berlin, Heidelberg. https://link.springer.com/content/pdf/10.1007%2F3-540-47721-7_12.pdf

¹⁵For all proof variants, ElectionGuard uses a compact representation that omits the commitments. Proofs consist of the challenge and response values only. When verifying proofs, the verification equations can be used to recompute the commitments and check their correctness via the challenge hash computation.

¹⁶Shamir A. (1979) *How to Share a Secret*. *Communications of the ACM*, vol 22.

use to form the partial decryptions. A pre-determined threshold quorum value (k) out of the (n) guardians will be necessary to produce a full decryption.

Encryption of other data

ElectionGuard provides means to encrypt data other than votes, which are selections usually encoded as zero or one that need to be homomorphically aggregated. Such data may include the cryptographic shares of a guardian’s secret key that are sent to the other guardians and are encrypted to the receiving guardians’ public keys, write-in information that needs to be attached to an encrypted selection and is encrypted to the election public key, or other auxiliary data attached to an encrypted ballot, either encrypted to the election public key or to an administrative public key. The non-vote data do not need to be homomorphically encrypted and can use a more standard form of public-key encryption removing the data size restrictions imposed by the vote encryption. ElectionGuard encrypts such data with hashed ElGamal encryption, which deploys a key derivation function (KDF) to generate a key stream that is then XORed with the data. To implement the KDF and to provide a message authentication code (MAC), encryption makes use of the keyed Hash Message Authentication Code HMAC. In ElectionGuard, HMAC is instantiated as HMAC-SHA-256 with the hash function SHA-256.

3.1 Parameter requirements

ElectionGuard uses integers to instantiate the encryption rather than elliptic curves in order to make construction of election verifiers as simple as possible without requiring special tools and dependencies. The encryption scheme used to encrypt votes is defined by a prime p and a prime q which divides $(p - 1)$. We use $r = (p - 1)/q$ to denote the cofactor of q , and a generator g of the order q subgroup \mathbb{Z}_p^* is fixed. We also require $r/2$ to be prime. The specific values for a 4096-bit prime p and a 256-bit prime q which divides $(p - 1)$ and a generator g that define all standard baseline parameters are given below.

3.1.1 Standard baseline cryptographic parameters

Standard parameters for ElectionGuard begin with the largest 256-bit prime $q = 2^{256} - 189$. The (big endian) hexadecimal representation of q is as follows.

$$q = \text{0xFFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF43} \quad (3)$$

The modulus p is then set to be a 4096-bit prime with the following properties.

1. The first 256 bits of p are all ones.
2. The last 256 bits of p are all ones.
3. $p - 1$ is a multiple of q .
4. $(p - 1)/2q$ is also prime.

The middle 3584 bits of p are chosen by starting with the first 3584 bits of the constant $\ln(2)$ (the natural logarithm of 2).¹⁷ After pre-pending and appending 256 ones, p is determined by finding the smallest prime larger than this value that satisfies the above properties.

¹⁷See <https://oeis.org/A068426> for the integer sequence consisting of the bits of $\ln(2)$.

This works out to $p = 2^{4096} - 2^{3840} + 2^{256}(\lfloor 2^{3584} \ln(2) \rfloor + \delta) + (2^{256} - 1)$ where the value of δ is given by

$$\delta = 287975203778583638958138611533602521491887169409704874643524560756486080635197037903.^{18}$$

The hexadecimal representation of p is as follows.

```
p = 0xFFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF
    B17217F7 D1CF79AB C9E3B398 03F2F6AF 40F34326 7298B62D 8A0D175B 8BAAFA2B
    E7B87620 6DEBAC98 559552FB 4AFA1B10 ED2EAE35 C1382144 27573B29 1169B825
    3E96CA16 224AE8C5 1ACBDA11 317C387E B9EA9BC3 B136603B 256FA0EC 7657F74B
    72CE87B1 9D6548CA F5DFA6BD 38303248 655FA187 2F20E3A2 DA2D97C5 0F3FD5C6
    07F4CA11 FB5BFB90 610D30F8 8FE551A2 EE569D6D FC1EFA15 7D2E23DE 1400B396
    17460775 DB8990E5 C943E732 B479CD33 CCCC4E65 9393514C 4C1A1E0B D1D6095D
    25669B33 3564A337 6A9C7F8A 5E148E82 074DB601 5CFE7AA3 0C480A54 17350D2C
    955D5179 B1E17B9D AE313CDB 6C606CB1 078F735D 1B2DB31B 5F50B518 5064C18B
    4D162DB3 B365853D 7598A195 1AE273EE 5570B6C6 8F969834 96D4E6D3 30AF889B
    44A02554 731CDC8E A17293D1 228A4EF9 8D6F5177 FBCF0755 268A5C1F 9538B982
    61AFFD44 6B1CA3CF 5E9222B8 8C66D3C5 422183ED C9942109 0BBB16FA F3D949F2
    36E02B20 CEE886B9 05C128D5 3D0BD2F9 62136319 6AF50302 0060E499 08391A0C
    57339BA2 BEBA7D05 2AC5B61C C4E9207C EF2F0CE2 D7373958 D7622658 90445744
    FB5F2DA4 B7510058 92D35689 0DEFE9CA D9B9D4B7 13E06162 A2D8FDD0 DF2FD608
    FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF
```

The hexadecimal representation of the cofactor $r = (p - 1)/q$ is as follows.

```
r = 0x01 00000000 00000000 00000000 00000000 00000000 00000000 00000000 000000BC
    B17217F7 D1CF79AB C9E3B398 03F2F6AF 40F34326 7298B62D 8A0D175B 8BAB857A
    E8F42816 5418806C 62B0EA36 355A3A73 E0C74198 5BF6A0E3 130179BF 2F0B43E3
    3AD86292 3861B8C9 F768C416 9519600B AD06093F 964B27E0 2D868312 31A9160D
    E48F4DA5 3D8AB5E6 9E386B69 4BEC1AE7 22D47579 249D5424 767C5C33 B9151E07
    C5C11D10 6AC446D3 30B47DB5 9D352E47 A53157DE 04461900 F6FE360D B897DF53
    16D87C94 AE71DAD0 BE84B647 C4BCF818 C23A2D4E BB53C702 A5C8062D 19F5E9B5
    033A94F7 FF732F54 12971286 9D97B8C9 6C412921 A9D86797 70F499A0 41C297CF
    F79D4C91 49EB6CAF 67B9EA3D C563D965 F3AAD137 7FF22DE9 C3E62068 DDOED615
    1C37B4F7 4634C2BD 09DA912F D599F433 3A8D2CC0 05627DCA 37BAD43E 64A39631
    19C0BFE3 4810A21E E7CFC421 D53398CB C7A95B3B F585E5A0 4B790E2F E1FE9BC2
    64FDA810 9F6454A0 82F5EFB2 F37EA237 AA29DF32 0D6EA860 C41A9054 CCD24876
    C6253F66 7BFB0139 B5531FF3 01899612 02FD2B0D 55A75272 C7FD7334 3F7899BC
```

¹⁸Discovering this value δ required enumerating roughly 2.49 million values satisfying the first three of the above properties to find the first one for which both p and $(p - 1)/2q$ are both prime.

A0B36A4C 470A64A0 09244C84 E77CEBC9 2417D5BB 13BF1816 7D8033EB 6C4DD787
9FD4A7F5 29FD4A7F 529FD4A7 F529FD4A 7F529FD4 A7F529FD 4A7F529F D4A7F52A

Finally, the generator g is chosen to be $g = 2^r \bmod p$ and has the following hexadecimal representation.

$g =$ 0x36036FED 214F3B50 DC566D3A 312FE413 1FEE1C2B CE6D02EA 39B477AC 05F7F885
F38CFE77 A7E45ACF 4029114C 4D7A9BFE 058BF2F9 95D2479D 3DDA618F FD910D3C
4236AB2C FDD783A5 016F7465 CF59BBF4 5D24A22F 130F2D04 FE93B2D5 8BB9C1D1
D27FC9A1 7D2AF49A 779F3FFB DCA22900 C14202EE 6C996160 34BE35CB CDD3E7BB
7996ADFE 534B63CC A41E21FF 5DC778EB B1B86C53 BFBE9998 7D7AEA07 56237FB4
0922139F 90A62F2A A8D9AD34 DFF799E3 3C857A64 68D001AC F3B681DB 87DC4242
755E2AC5 A5027DB8 1984F033 C4D17837 1F273DBB 4FCEA1E6 28C23E52 759BC776
5728035C EA26B44C 49A65666 889820A4 5C33DD37 EA4A1D00 CB62305C D541BE1E
8A92685A 07012B1A 20A746C3 591A2DB3 815000D2 AACCFE43 DC49E828 C1ED7387
466AFD8E 4BF19355 93B2A442 EEC271C5 0AD39F73 3797A1EA 11802A25 57916534
662A6B7E 9A9E449A 24C8CFF8 09E79A4D 806EB681 119330E6 C57985E3 9B200B48
93639FDF DEA49F76 AD1ACD99 7EBA1365 7541E79E C57437E5 04EDA9DD 01106151
6C643FB3 0D6D58AF CCD28B73 FEDA29EC 12B01A5E B86399A5 93A9D5F4 50DE39CB
92962C5E C6925348 DB54D128 FD99C14B 457F883E C20112A7 5A6A0581 D3D80A3B
4EF09EC8 6F9552FF DA1653F1 33AA2534 983A6F31 BOEE4697 935A6B1E A2F75B85
E7EBA151 BA486094 D68722B0 54633FEC 51CA3F29 B31E77E3 17B178B6 B9D8AE0F

Alternative parameter sets are possible and may be allowed in future versions of ElectionGuard¹⁹.

Note 3.1. As an example for alternative parameters, the Appendix provides a set of reduced parameters that offer better performance at a lower security level, and additionally provides various sets of very small and insecure parameters for testing purposes only. A good source for parameter generation is appendix A of FIPS 186-4²⁰. Allowing alternate non-standard parameters would force election verifiers to recognize and check that parameters are correctly generated. Since these checks are very different from other checks that are required of a verifier, allowing such parameters would add substantial complexity to election verifiers. For this reason, this version of ElectionGuard fixes the parameters as above.

¹⁹If alternative parameters are allowed, election verifiers must confirm that p , q , r , and g are such that both p and q are prime (this may be done probabilistically using the Miller-Rabin algorithm), that $p - 1 = qr$ is satisfied, that q is not a divisor of r , that $1 < g < p$, that $g^q \bmod p = 1$, and that generation of the parameters is consistent with the cited standard.

²⁰NIST (2013) Digital Signature Standard (DSS). In: FIPS 186-4. <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf>

3.1.2 Parameter base hash

The prime modulus p , the subgroup order q , and the subgroup generator g are hashed using the ElectionGuard hash function H (as specified in detail in Section 5) to form the parameter base hash

$$H_P = H(\text{ver}; 0x00, p, q, g). \quad (4)$$

The symbol `ver` denotes the version byte array that encodes the used version of this specification. The array has length 32 and contains the UTF-8 encoding of the string “v2.0.0” followed by 0x00-bytes, i.e. `ver = 0x76322E302E30 || b(0, 27)`. The 0x00-byte at the beginning of the second argument of H is a domain separation byte (written in hexadecimal notation), see Section 5.1. For the baseline parameters specified in Section 3.1.1 above, the hexadecimal representation of the parameter base hash is

$$H_P = 0x2B3B025E50E09C119CBA7E9448ACD1CABC9447EF39BF06327D81C665CDD86296. \quad (5)$$

This hash value is an input to all subsequent hash computations, either directly or indirectly, which binds all hash outputs to the baseline parameters and the version of the specification used for the election.

3.1.3 Election parameters and the election manifest

Another parameter of an election is a public election manifest. The ElectionGuard manifest is a single file that describes the modalities of the election and must specify a wide range of data about it. Most importantly, the manifest lists, among other information, all the contests in an election, the candidates (selectable options) for each contest together with associations between each option and its representation on a virtual ballot, and the contest selection limit—the number of options that a voter may select in that contest. It is assumed that each contest in the manifest has a unique label and that within each contest, each option also has a unique label. For instance, if Alice, Bob, and Carol are running for governor, and David and Ellen are running for senator, the election manifest file could enable the vector $\langle 0, 1, 0; 0, 1 \rangle$ to be recognized as a ballot with votes for Bob as governor and Ellen as senator. This section defines how the manifest must uniquely specify the election parameters.

Labels. The election objects relevant for ElectionGuard such as contests, selectable options, and ballot styles are identified by unique labels. A *label* is a string, it should be a short, concise name or identifier and should therefore not contain any long-form descriptions or line break characters, tabs, and similar special characters. Likewise, a label should not have leading or trailing whitespace characters.

The ElectionGuard manifest must contain a label that contains a unique and descriptive identifier for the election.

Indices. Instead of using the labels directly, ElectionGuard handles contests, selectable options, ballot styles, etc. in terms of their position in a unique ordered list. They can therefore be uniquely

identified by an index value. An *index* is a 1-based ordinal in the range $1 \leq i < 2^{31}$.²¹

Contests and the contest index. The election manifest must contain a single ordered list of all the contests that can appear on any ballot generated for the election. Each contest must have a contest label $\lambda_{\mathcal{C}}$ that is unique within the election. The position of a given contest \mathcal{C} in this list is the *contest index* $\text{ind}_{\mathcal{C}}(\lambda_{\mathcal{C}})$. The contest list provides a bijective mapping between the unique contest labels and the corresponding contest index values. Note that the first contest in the list has contest index 1.

Selectable options and the option index. For each contest in the contest list, the election manifest must also contain a single ordered list of all the selectable options in this contest. Each selectable option \mathcal{O} must have an option label $\lambda_{\mathcal{O}}$ that is unique within the contest. The position of a given selectable option \mathcal{O} in this list is the *option index* $\text{ind}_{\mathcal{O}}(\lambda_{\mathcal{O}})$. This option list provides a bijective mapping between the unique selectable option labels and the corresponding option index values. The first selectable option in the list has option index 1.

Ballot styles and the ballot style index. A *ballot style* is a single list of contest indices that defines which contests appear on a ballot generated according to this ballot style. If a contest index is contained in this list, the corresponding contest is present, otherwise it is not. The list is not ordered and only specifies whether a contest is present or not on a ballot of this ballot style. Each ballot style \mathcal{S} must have a ballot style label $\lambda_{\mathcal{S}}$ that is unique within the election. The election manifest must contain a single, ordered list of all ballot styles that are possible in the election. The position of a given ballot style \mathcal{S} in this list is the *ballot style index* $\text{ind}_{\mathcal{S}}(\lambda_{\mathcal{S}})$. The ballot style list provides a bijective mapping between the unique ballot styles and the corresponding ballot style index values. Again, the first ballot style in the list has ballot style index 1.

Selections, option selection limits and contest selection limits. A *selection* is the assignment of a value to an option by the voter. This could be the value 1, meaning that the voter selected this option, or the value 0, meaning that the voter did not select this option. To allow other voting methods, ElectionGuard allows a selection to be the assignment of a value in a range $\{0, 1, \dots, R\}$, where R is the *option selection limit*, a positive integer that defines the maximal value allowed to be assigned to this option by the voter. For each contest, the election manifest must specify the option selection limit for the options in this contest, and must also specify a *contest selection limit* L , which is the maximal total value for the sum of all selections made in that contest.

Optional data fields. The election manifest offers optional data fields for additional data that can be attached to any of the above concepts and to the manifest in general at the top level. They can be used to provide additional information on contests, selectable options, and ballot styles that go beyond the short unique labels described above. In addition, there can be data fields that describe general information about the election, jurisdiction, election device manufacturers, software versions, the date and location, etc. that help to interpret the additional data.

²¹The upper bound of $2^{31} - 1$ is in consideration of languages and runtimes that do not have full support for unsigned integers.

There are many other things that a manifest may specify. Some examples are listed here.

Undervotes. An undervote occurs in a contest if the number of selected options (or more generally, the total sum of selections assigned by the voter) is strictly less than the contest selection limit, the number (or total sum) of allowed selections a voter can make in that contest.

Counting undervoted contests. A manifest may specify whether the fact that a contest was undervoted should be recorded for public verification. If this is specified, the indicated contest should have an additional field that functions very much like a selectable option field. The value of this field is an encryption of one if the contest was undervoted on the ballot and an encryption of zero otherwise. The undervote field should be set to one only when the number of option fields that are set to one is strictly less than the contest selection limit. This can be enforced by checking that the sum of all field contents should not exceed the contest selection limit for that contest.

This field is an undervote indicator field and is added to all ballots that contain this contest if the occurrence of an undervote is to be verifiably recorded. The field only indicates that an undervote has occurred on the specific ballot in this contest. A tally of this field will contain the total number of ballots that showed an undervote in this contest. The field does not record by how many votes the contest was undervoted, i.e. by how much the total number of selections by the voter is less than the contest selection limit. This number can be recorded in an undervote difference count field.

Undervote Difference Count. A manifest may specify whether the total number of undervotes, i.e. the difference between the number of selections the voter made and the contest selection limit, for each contest should be recorded for public verification. If this is specified, the indicated contest should have an additional field that functions very much like an option field but whose value need not be limited to zero or one. The value of this field is an encryption of the number of undervotes on that ballot for that contest. This can be enforced by checking that the sum of all field contents should exactly match the contest selection limit for that contest.

Overvotes. An overvote occurs in a contest if the voter selected more options than allowed, i.e. more than the contest selection limit for that contest specifies. A manifest may specify whether an overvoted contest should be recorded for public verification. If this is specified, the indicated contest should have an additional field that functions very much like an option field. The value of this field is an encryption of one if the contest was overvoted on the ballot and an encryption of zero otherwise. When the overvote field is set to one, all other selection fields should be set to zero, and this can be enforced by checking that the sum of the other fields—when added to the product of the contest selection limit with the overvote field contents should not exceed the selection limit for that contest.²² Similar to the undervote difference count, the overvote difference count can be verifiably recorded in a separate overvote difference field.

²²An alternative implementation would be to always set the overvote field to either zero or the contest selection limit. This would be slightly more efficient computationally, but the publicly verified number of overvotes for that contest would be scaled up by a factor of the contest selection limit and the selection limit would then need to be divided out for reporting purposes.

Write-Ins. A manifest may specify whether the utilization of each write-in field in each contest is recorded for public verification. If this is specified, the indicated contest should have one or more additional fields (usually matching the contest selection limit) that function very much like option fields. If a single write-in is utilized in a contest, the first write-in field is set to one and all others are set to zero; if two write-ins are utilized in a contest, the first two write-in fields are set to one and all others are set to zero; etc. Range proofs should be included to show that each of these individual write-in fields is an encryption of either zero or one.

Write-Ins Total. A manifest may specify whether the total number of write-ins selected in each contest is recorded for public verification. If this is specified, the indicated contest should have a single additional field that functions very much like option fields but whose value need not be limited to zero or one. The value of this field is an encryption of the number of write-ins utilized on that ballot for that contest. Range proofs should be used to demonstrate that in each instance, the value of this field is non-negative and does not exceed the contests selection limit.

A note on selection limit proofs. In general, a combination of the above fields can be merged into a single selection limit proof per contest per ballot. For example, a contest which includes an undervote field, an overvote field, and two separate write-in fields can include a single range proof that the sum of these fields is non-negative and does not exceed the contest selection limit. The only exceptions are the very unusual cases where a manifest specifies including both an undervote field and an undervote count field or both individual write-in fields and a write-in total field. In these cases, the selection limit proofs should each include only one of the paired fields (e.g., undervote or undervote count) and a second selection limit proof should be included if both paired fields are to be verified.

3.1.4 Election manifest hash

The data in the election manifest is written to a file `manifest` in a canonical representation that may be implementation specific. The election manifest file `manifest` is hashed²³ using the ElectionGuard hash function with the parameter hash H_P to produce the manifest digest

$$H_M = H(H_P; 0x01, \text{manifest}). \quad (6)$$

3.1.5 Election base hash

The election base hash H_B is computed from the parameter base hash H_P and the manifest hash H_M together with the number n of guardians that participate in the election, and the quorum value k , as

$$H_B = H(H_P; 0x02, H_M, n, k). \quad (7)$$

Incorporating the byte array H_B into subsequent hash computations binds those hashes to the election parameters, the number of guardians, the threshold value and to a specific election by including the manifest.

²³The file that constitutes the election manifest is input to H as an entire file. Again the `0x01` byte at the beginning of the second argument is a domain separation byte (in hex). In what follows, all uses of the function H include such domain separation bytes.

Verification 1 (Parameter validation)

An ElectionGuard election verifier must verify that it uses the correct version of the ElectionGuard specification, that it uses the standard baseline parameters, which may be hardcoded, and that the base hash values have been computed correctly.

- (1.A) The ElectionGuard specification version used to generate the election record is the same as the ElectionGuard specification version used to verify the election record.
- (1.B) The large prime is equal to the large modulus p defined in Section 3.1.1.
- (1.C) The small prime is equal to the prime q defined in Section 3.1.1.
- (1.D) The cofactor is equal to the value r defined in Section 3.1.1.
- (1.E) The generator is equal to the generator g defined in Section 3.1.1.
- (1.F) The parameter base hash has been computed correctly as

$$H_P = H(\text{ver}; 0x00, p, q, g)$$

using the version byte array $\text{ver} = 0x76322E302E30 \parallel \mathbf{b}(0, 26)$, which is the UTF-8 encoding of the version string “v2.0.0” padded with 0x00-bytes to length 32 bytes.

- (1.G) The manifest hash has been computed correctly from the manifest as

$$H_M = H(H_P; 0x01, \text{manifest}).$$

- (1.H) The base hash has been computed correctly as

$$H_B = (H_P; 0x02, H_M, n, k).$$

3.2 Key Generation

Before an election, the number of guardians (n) is fixed together with a quorum value (k) that describes the number of guardians necessary to decrypt tallies and produce election verification data. The values n and k are integers subject to the constraint that $1 \leq k \leq n$. Canvassing board members can often serve the role of election guardians, and typical values for n and k could be 5 and 3 – indicating that 3 of 5 canvassing board members must cooperate to produce the artifacts that enable election verification. The reason for not setting the quorum value k too low is that it will also be possible for k guardians to decrypt individual ballots.

Note 3.2. Decryption of individual ballots does not directly compromise voter privacy since links between encrypted ballots and the voters who cast them are not retained by the system. However, voters receive confirmation codes that can be associated with individual encrypted ballots, so any group that has the ability to decrypt individual ballots can also coerce voters by demanding to see their confirmation codes.

Threshold encryption is used for encryption of ballots. This form of encryption makes it very easy to combine individual guardian public keys into a single public key for encrypting votes and ballots. It also offers a homomorphic property that allows individual encrypted votes to be combined to form encrypted tallies.

The guardians of an election will each generate a public-secret key pair. The public keys will then be combined (as described in the following section) into a single election public key which is used to encrypt all selections made by voters in the election.

At the conclusion of the election, each guardian will compute a verifiable partial decryption of each tally. These partial decryptions will then be combined to form full verifiable decryptions of the election tallies.

To accommodate the possibility that one or more of the guardians will not be available at the conclusion of the election to form their partial decryptions, the guardians will cryptographically share²⁴ their secret keys amongst each other during key generation in a manner to be detailed in the next section. Each guardian will then compute a share of the secret decryption key, which it uses to form the partial decryptions. A pre-determined quorum value (k) out of the (n) guardians will be necessary to produce a full decryption.

If the same set of n guardians support multiple elections using the same threshold value k , the generated keys and key shares may be reused across several elections.

3.2.1 Overview of key generation

The n guardians of an election are denoted by G_1, G_2, \dots, G_n . Each guardian G_i generates an independent public-secret key pair by generating a random integer secret $s_i \in \mathbb{Z}_q$ and forming the public key $K_i = g^{s_i} \bmod p$. Each of these public keys will be published in the election record together with a non-interactive zero-knowledge Schnorr proof of knowledge of the associated secret key.

The joint election public key will be

$$K = \prod_{i=1}^n K_i \bmod p. \quad (8)$$

To enable robustness and allow for the possibility of missing guardians at the conclusion of an election, the ElectionGuard key generation includes a sharing of secret keys between guardians to enable decryption by any k guardians. This sharing is verifiable, so that receiving guardians can confirm that the shares they receive are correct; and the process allows for decryption without explicitly reconstructing secret keys of missing guardians.

Each guardian G_i generates $k - 1$ random polynomial coefficients $a_{i,j}$ such that $0 < j < k$ and $0 \leq a_{i,j} < q$ and forms the polynomial

$$P_i(x) = \sum_{j=0}^{k-1} a_{i,j} x^j \bmod q$$

by setting $a_{i,0}$ equal to its secret value s_i . Guardian G_i then publishes commitments $K_{i,j} = g^{a_{i,j}} \bmod p$ to each of its polynomial coefficients. As with the primary secret keys, each guardian should provide a Schnorr proof of knowledge of the secret coefficient value $a_{i,j}$ associated with each published commitment $K_{i,j}$. Since polynomial coefficients will be generated and managed in precisely the same fashion as secret keys, they will be treated together in a single step below.

²⁴Shamir A. (1979) *How to Share a Secret*. Communications of the ACM, vol 22.

At the conclusion of the election, individual encrypted ballots will be homomorphically combined into a single encrypted aggregate ballot – consisting of an encryption of the tally for each option offered to voters. Each guardian will use its share of the secret decryption key to generate a partial decryption of each encrypted tally value, and these partial decryptions will be combined into full decryptions. If any election guardians are missing during tallying, any set of k guardians who are available can cooperate to complete decryption.

All challenged ballots are individually decrypted in precisely the same fashion.

3.2.2 Details of key generation

Guardian coefficients and key pair. Each guardian G_i in an election with a decryption threshold of k generates k secret polynomial coefficients $a_{i,j}$, for $0 \leq j < k$, by sampling them uniformly, at random in the range $0 < a_{i,j} < q$ and forms the polynomial

$$P_i(x) = \sum_{j=0}^{k-1} a_{i,j} x^j \text{ mod } q. \quad (9)$$

Guardian G_i then publishes commitments

$$K_{i,j} = g^{a_{i,j}} \text{ mod } p \quad (10)$$

for each of its polynomial coefficients $a_{i,j}$, for $0 \leq j < k$. The constant term $a_{i,0}$ of this polynomial will serve as the secret key for guardian G_i , and for convenience we denote $s_i = a_{i,0}$, and its commitment $K_{i,0}$ will serve as the public key for guardian G_i and will also be denoted as $K_i = K_{i,0}$. In order to prove possession of the coefficient associated with each public commitment, for each $K_{i,j}$ with $0 \leq j < k$, guardian G_i generates a Schnorr proof of knowledge for each of its coefficients as follows.

NIZK Proof: Guardian G_i proves knowledge of secrets $a_{i,j}$ such that $K_{i,j} = g^{a_{i,j}} \text{ mod } p$.

For each $0 \leq j < k$, Guardian G_i generates random integer values $u_{i,j}$ in \mathbb{Z}_q and computes

$$h_{i,j} = g^{u_{i,j}} \text{ mod } p. \quad (11)$$

Then, using the hash function H (as defined in Section 5 below), guardian G_i performs a single hash computation

$$c_{i,j} = H(\text{HP}; 0x10, i, j, K_{i,j}, h_{i,j}) \quad (12)$$

and publishes the Schnorr proof $(c_{i,j}, v_{i,j})$ consisting of the challenge value $c_{i,j}$ and the response value $v_{i,j} = (u_{i,j} - c_{i,j} a_{i,j}) \text{ mod } q$ together with $K_{i,j}$.

Verification 2 (Guardian public-key validation)

For each guardian G_i , $1 \leq i \leq n$, and for each $j \in \mathbb{Z}_k$, an election verifier must compute the value

$$(2.1) \quad h_{i,j} = g^{v_{i,j}} \cdot K_{i,j}^{c_{i,j}} \pmod p$$

and then must confirm the following.

- (2.A) The value $K_{i,j}$ is in \mathbb{Z}_p^r . (A value x is in \mathbb{Z}_p^r if and only if x is an integer such that $0 \leq x < p$ and $x^q \pmod p = 1$ is satisfied.)
- (2.B) The value $v_{i,j}$ is in \mathbb{Z}_q . (A value x is in \mathbb{Z}_q if and only if x is an integer such that $0 \leq x < q$.)
- (2.C) The challenge $c_{i,j}$ is correctly computed as $c_{i,j} = H(\text{HP}; 0\text{x}10, i, j, K_{i,j}, h_{i,j})$.

25

It is worth noting here that for any fixed constant α , the value $g^{P_i(\alpha)} \pmod p$ can be computed entirely from the published commitments as

$$g^{P_i(\alpha)} \pmod p = g^{\sum_{j=0}^{k-1} a_{i,j} \alpha^j} \pmod p = \prod_{j=0}^{k-1} g^{a_{i,j} \alpha^j} \pmod p = \prod_{j=0}^{k-1} (g^{a_{i,j}})^{\alpha^j} \pmod p = \prod_{j=0}^{k-1} K_{i,j}^{\alpha^j} \pmod p. \quad (13)$$

Note 3.3. Although this formula includes double exponentiation – raising a given value to the power α^j – in what follows, α and j will always be small values (bounded by n). This can also be reduced since the same result will be achieved if the exponents α^j are reduced to $\alpha^j \pmod q$.

Share encryption. To share secret values amongst each other, each guardian G_i encrypts its share for each other guardian G_ℓ using an encryption function E_ℓ that can be instantiated using that guardian’s public/secret key pair ($K_\ell = g^{s_\ell} \pmod p, s_\ell$) as laid out in the section describing the encryption of other data above. This means that each guardian G_i publishes the encryption $E_\ell(P_i(\ell))$ for every other guardian G_ℓ as follows.

Guardian G_i uses guardian G_ℓ ’s public key K_ℓ and a randomly-selected nonce $\xi_{i,\ell} \in \mathbb{Z}_q$ to compute

$$(\alpha_{i,\ell}, \beta_{i,\ell}) = (g^{\xi_{i,\ell}} \pmod p, K_\ell^{\xi_{i,\ell}} \pmod p) \quad (14)$$

and the secret key

$$k_{i,\ell} = H(\text{HP}; 0\text{x}11, i, \ell, K_\ell, \alpha_{i,\ell}, \beta_{i,\ell}). \quad (15)$$

Using the keyed Hash Message Authentication Code HMAC instantiated with SHA-256, this key is used to derive²⁶ the MAC key

$$k_0 = \text{HMAC}(k_{i,\ell}, 0\text{x}01 \parallel \text{Label} \parallel 0\text{x}00 \parallel \text{Context} \parallel 0\text{x}0200) \quad (16)$$

and the encryption key

$$k_1 = \text{HMAC}(k_{i,\ell}, 0\text{x}02 \parallel \text{Label} \parallel 0\text{x}00 \parallel \text{Context} \parallel 0\text{x}0200), \quad (17)$$

²⁵This verification box and some others below contain computation steps as well as verification conditions. The former are numbered with decimal numerals, while the latter are numbered with capital letters alphabetically.

²⁶NIST (2022) *Recommendation for Key Derivation Using Pseudorandom Functions*. In: SP 800-108r1 <https://csrc.nist.gov/publications/detail/sp/800-108/rev-1/final>.

which both are 256 bits (32 bytes) long, and where `Label` = `b(“share_enc_keys”, 14)` and `Context` = `b(“share_encrypt”, 13) || b(i, 4) || b(l, 4)`.²⁷

Since $P_i(\ell)$ is an integer modulo q , it can be encoded as a byte array `b(Pi(ℓ), 32)` of exactly 32 bytes as specified in Section 5.1.2. The encoded value is then encrypted as

$$E_\ell(P_i(\ell)) = (C_{i,\ell,0}, C_{i,\ell,1}, C_{i,\ell,2}), \quad (18)$$

where

$$C_{i,\ell,0} = g^{\xi_{i,\ell}} \bmod p, \quad C_{i,\ell,1} = \mathbf{b}(P_i(\ell), 32) \oplus k_1, \quad C_{i,\ell,2} = \text{HMAC}(k_0, \mathbf{b}(C_{i,\ell,0}, 512) \parallel C_{i,\ell,1}). \quad (19)$$

Share decryption. After receiving the ciphertext $(C_{i,\ell,0}, C_{i,\ell,1}, C_{i,\ell,2})$ from guardian G_i , guardian G_ℓ decrypts it by computing $\beta_{i,\ell} = (C_{i,\ell,0})^{s_\ell} \bmod p$, setting $\alpha_{i,\ell} = C_{i,\ell,0}$ and obtaining $k_{i,\ell} = H(\text{HP}; \text{0x11}, i, \ell, K_\ell, \alpha_{i,\ell}, \beta_{i,\ell})$. Now the MAC key k_0 and the encryption key k_1 can be computed as above in Equations (16) and (17), which allows G_ℓ to verify the validity of the MAC, namely that $C_{i,\ell,2} = \text{HMAC}(k_0, \mathbf{b}(C_{i,\ell,0}, 512) \parallel C_{i,\ell,1})$. If the MAC verifies, G_ℓ decrypts

$$\mathbf{b}(P_i(\ell), 32) = C_{i,\ell,1} \oplus k_1. \quad (20)$$

Share verification. Having decrypted each $P_i(\ell)$, guardian G_ℓ can now verify its validity against the commitments $K_{i,0}, K_{i,1}, \dots, K_{i,k-1}$ made by G_i to its coefficients by confirming that the following equation holds:

$$g^{P_i(\ell)} \bmod p = \prod_{j=0}^{k-1} (K_{i,j})^{\ell^j} \bmod p. \quad (21)$$

Guardians then publicly report having confirmed or failed to confirm this computation. If the recipient guardian G_ℓ reports not receiving a suitable value $P_i(\ell)$, it becomes incumbent on the sending guardian G_i to publish this $P_i(\ell)$ together with the nonce $\xi_{i,\ell}$ it used to encrypt $P_i(\ell)$ under the public key K_ℓ of recipient guardian G_ℓ . If guardian G_i fails to produce a suitable $P_i(\ell)$ and nonce $\xi_{i,\ell}$ that match both the published encryption and the above equation, it should be excluded from the election and the key generation process should be restarted with an alternate guardian. If, however, the published $P_i(\ell)$ and $\xi_{i,\ell}$ satisfy both the published encryption and the equation above, the claim of malfeasance is dismissed, and the key generation process continues undeterred.²⁸

Verification 3 (Election public-key validation)

An election verifier must verify the correct computation of the joint election public key.

(3.A) The value K_i is in \mathbb{Z}_p^r and $K_i \neq 1 \bmod p$ for all $1 \leq i \leq n$.

(3.B) $K = \prod_{i=1}^n K_i \bmod p$ and $K \neq 1 \bmod p$.

²⁷This key derivation uses the KDF in counter mode from SP 800-108r1. The second input to HMAC contains the counter in the first byte, the UTF-8 encoding of the string “share_enc_keys” as the *Label* (encoding is denoted by `b(...)`, see Section 5.1.4), a separation `0x00` byte, the UTF-8 encoding of the string “share_encrypt” concatenated with encodings of the numbers i and ℓ of the sending and receiving guardians as the *Context*, and the final two bytes specifying the length of the output key material as 512 bits in total.

²⁸It is also permissible to dismiss any guardian that makes a false claim of malfeasance. However, this is not required as the sensitive information that is released as a result of the claim could have been released by the claimant in any case.

Each guardian G_i will also need to compute the value

$$P(i) = \left(\sum_{j=1}^n P_j(i) \right) \bmod q = (P_1(i) + P_2(i) + \dots + P_n(i)) \bmod q. \quad (22)$$

for later use in decryption as will be described in Section 3.6.1.

At a minimum, guardian G_i should retain the values s_i and $P(i)$.

3.2.3 Extended base hash

Once the baseline parameters have been produced and confirmed, the base hash H_B is hashed with the election public key K to form an extended base hash

$$H_E = H(H_B; 0x12, K) \quad (23)$$

that will form the basis of subsequent hash computations.

Verification 4 (Extended base hash validation)

An election verifier must verify the correct computation of the extended base hash.

(4.A) $H_E = H(H_B; 0x12, K)$.

3.3 Ballot Encryption

Although there are some exceptions, an ElectionGuard ballot is typically comprised entirely of encryptions of one (indicating selection made) and zero (indicating selection not made). To enable homomorphic addition (for tallying), these values are exponentiated during vote encryption.

3.3.1 Selection encryption

To encrypt a ballot entry σ , a random value ξ is selected such that $0 \leq \xi < q$, and σ is encrypted as

$$\text{Enc}(\sigma, \xi) = (\alpha, \beta) = (g^\xi \bmod p, K^\sigma \cdot K^\xi \bmod p) = (g^\xi \bmod p, K^{\sigma+\xi} \bmod p). \quad (24)$$

Specifically, for the typical case of $\sigma \in \{0, 1\}$, one of the following two computations is performed.

- Zero (not selected, $\sigma = 0$) is encrypted as $\text{Enc}(0, \xi) = (g^\xi \bmod p, K^\xi \bmod p)$.
- One (selected, $\sigma = 1$) is encrypted as $\text{Enc}(1, \xi) = (g^\xi \bmod p, K^{1+\xi} \bmod p)$.

The notation $\text{Enc}(\sigma, \xi)$ for this homomorphically additive encryption includes the encryption nonce ξ . In cases, where the nonce does not need to be referenced, $\text{Enc}(\sigma)$ is sometimes used for clarity.

Note that if multiple encrypted votes $(g^{\xi_i} \bmod p, K^{\sigma_i+\xi_i} \bmod p)$ are formed, their component-wise product $(g^{\sum_i \xi_i} \bmod p, K^{\sum_i \sigma_i + \sum_i \xi_i} \bmod p)$ serves as an encryption of $\sum_i \sigma_i$ —which is the tally of those votes.²⁹

²⁹The initial decryption actually forms the value $g^{\sum_i \sigma_i} \bmod p$. However, since $\sum_i \sigma_i$ is a relatively small value, it can be effectively computed from $g^{\sum_i \sigma_i} \bmod p$ by means of an exhaustive search or similar methods.

Some cardinal voting methods like cumulative voting, score voting, STAR-voting, and Borda count may require the encryption of small positive integers σ greater than 1 that represent multiple allowed votes or weighted votes. In all these cases, encryption is done as shown in Equation (24).

3.3.2 Generation of encryption nonces

The “random” nonces used for vote encryption on a single ballot B are all derived from a single 256-bit *ballot nonce* ξ_B . It is assumed that each contest has a label Λ and that within each contest each possible option has a label λ , which are specified in the election manifest. The manifest also specifies a unique sequence order of contests and options within contests. Within the i -th contest, which has contest label Λ_i , the nonce $\xi_{i,j}$ used to encrypt the selection for the j -th option, which has option label λ_j , is derived as

$$\xi_{i,j} = H(H_E; 0x20, \xi_B, \text{ind}_c(\Lambda_i), \text{ind}_o(\lambda_j)). \quad (25)$$

Ballot nonces may be independent across different ballots, and only the nonces used to encrypt ballot selections need to be derived from the ballot nonce.³⁰ The use of a single ballot nonce for each ballot allows the entire ballot encryption to be re-derived from the contents of a ballot and the ballot nonce. It also allows the encrypted ballot to be fully decrypted with the single ballot nonce.

Since access to a ballot nonce ξ_B enables decryption of the ballot which it encrypted, ballot nonces must be protected carefully. Typically, ballot nonces will be destroyed immediately after they are used. But there are several scenarios which warrant the preservation of ballot nonces—sometimes only briefly, but sometimes for an extended period.

To accommodate preservation of ballot nonces, an ElectionGuard manifest may optionally provide a public key to be used for encryption of ballot nonces. This public key may be the same as the election key used to encrypt voter selections, but some scenarios warrant the use of an independent public key. If such a key is provided, ElectionGuard uses that key to encrypt each ballot’s ballot nonce ξ_B and return this encryption together with the encrypted ballot. There are also some scenarios in which ballot nonces are encrypted with symmetric keys. Any such symmetric key should never be published, and the process for managing these symmetric keys is outside of the scope of this specification.

3.3.3 Ballot well-formedness

Contest selection limits. A contest in an election consists of a set of options together with a selection limit that indicates the number of selections that are allowed to be made in that contest. In most elections, most contests have a selection limit of one. However, a larger selection limit (e.g., select up to three) is not uncommon in some elections. Approval voting can be achieved by setting the selection limit to the total number of options in a contest. Ranked choice voting is not supported in this version of ElectionGuard, but it may be enabled in a future version.³¹ Write-ins

³⁰As will be seen below, in addition to encryption of voter selections, an ElectionGuard ballot contains additional encryptions that enable independent verification of a ballot’s correctness. The nonces used for these additional encryptions do not need to be derived from the ballot nonce ξ_B .

³¹Benaloh J., Moran. T, Naish L., Ramchen K., and Teague V. (2009) *Shuffle-Sum: Coercion-Resistant Verifiable Tallying for STV Voting* in Transactions of Information Forensics and Security.

are assumed to be explicitly registered or allowed to be lumped into a single “write-ins” category for the purpose of verifiable tallying. Verifiable tallying of free-form write-ins may be best done with a MixNet³² design.

Undervotes. A legitimate vote in a contest consists of a set of selections with cardinality not exceeding the selection limit of that contest. To accommodate legitimate undervotes, in previous versions of ElectionGuard (up to v1.1), the internal representation of a contest was augmented with “placeholder” options equal in number to the selection limit. Placeholder options were selected as necessary to force the total number of selections made in a contest to be equal to the selection limit. When the selection limit is one, for example, ElectionGuard used a single placeholder option that can be thought of as a “none of the above” option. The current version of ElectionGuard adopts a different approach to accommodate undervotes and allows the total number of selected options to lie in a range between 0 and the selection limit. This has lower computational cost and leads to smaller election records.

Overvotes. When the number of selections made by the voter exceeds the selection limit for a specific contest, the votes in the contest become invalid as an overvote. To not affect the election tallies, all selectable options in the contest are set to zero (unselected). To record that an overvote has occurred and what specific selections were made by the voter in this contest, this information is encoded into a contest data field, together with other data such as write-in text. This data is then encrypted with the election public key using the hashed ElGamal encryption described in Section 3.3.6.

Ballot well-formedness . Two things must now be proven about the encryption of each vote to ensure a ballot is well-formed.

1. The encryption associated with each option is an encryption of a legitimate value—usually either an encryption of zero or an encryption of one.
2. The sum of all encrypted selections in each contest lies in the range between 0 and the selection limit L for that contest (usually $L = 1$), i.e. it is an encryption of one of the values $0, 1, \dots, L$.

The use of DPP vote encryption enables efficient zero-knowledge proofs of these requirements, and the Fiat-Shamir heuristic can be used to make these proofs non-interactive. Chaum-Pedersen proofs are used to demonstrate that an encryption is that of a specified value, and these are combined with the Cramer-Damgård-Schoenmakers technique to show that an encryption is that of one of a specified set of values—usually that a value is an encryption of either zero or one. The encryptions of selections in a contest are homomorphically combined, and the result is shown to be an encryption of a non-negative value that is at most the contest’s selection limit—again using Chaum-Pedersen proofs combined with the Cramer-Damgård-Schoenmakers technique.

Cardinal voting methods. A range proof, i.e. a proof that a ciphertext is an encryption of one of the values $0, 1, \dots, L$ is a generalization of the proof that a ciphertext is an encryption of zero or one. It can be used to prove well-formedness of an individual selection when it is allowed for

³²Chaum D. (1981) *Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms* Communications of the ACM.

options to receive multiple votes in a contest by a single voter. Using range proofs with a range up to a certain option selection limit for the individual option as well as the contest selection sum therefore enables cardinal voting methods such as cumulative voting, score voting, STAR-voting, and Borda count in ElectionGuard.

3.3.4 Outline for proofs of ballot correctness

This section provides the outlines for different types of encryption proofs that aim to give an understanding of how they work. It starts with the proof that a given ciphertext is an encryption of the value 0. Based on that, it explains how it can be proved that a ciphertext is an encryption of the value 1. It then combines those to arrive at the base case used in ElectionGuard, namely that a ciphertext is an encryption of 0 or 1. Finally, it shows how the latter is generalized to generate range proofs, i.e. proofs that a ciphertext encrypts a value in a range $0, 1, \dots, L$.

NIZK Proof that (α, β) is an encryption of zero. To prove that ciphertext (α, β) is an encryption of zero, the Chaum-Pedersen protocol proceeds as follows. This proof assumes knowledge of the secret encryption nonce ξ . The prover selects a random value u in \mathbb{Z}_q and commits to the pair $(a, b) = (g^u \bmod p, K^u \bmod p)$. A hash computation is then performed (using the Fiat-Shamir heuristic) to create a pseudo-random challenge value $c = H(H_E, K, \alpha, \beta, a, b)$, and the prover responds with $v = (u - c\xi) \bmod q$. A verifier can now confirm the claim by checking that both $g^v \cdot \alpha^c \equiv_p a$ and $K^v \cdot \beta^c \equiv_p b$ are true.

NIZK Proof that (α, β) is an encryption of one. To prove that (α, β) is an encryption of one, $\beta/K \bmod p$ is substituted for β in the above. The verifier can be relieved of the need to perform a modular division by computing $\beta K^{q-1} \bmod p$ rather than $\beta/K \bmod p$. As an alternative, the verifier can confirm that $K^{(v-c) \bmod q} \cdot \beta^c \equiv_p b$ instead of $K^v \cdot (\beta/K)^c \equiv_p b$.

As with many zero-knowledge protocols, if the prover knows a challenge value prior to making its commitment, it can create a false proof. For example, if a particular challenge c is known to be forthcoming, a prover can generate a random $v \in \mathbb{Z}_q$ and commit to $(a, b) = (g^v \alpha^c \bmod p, K^v \beta^c \bmod p)$. This selection will satisfy the required checks for (α, β) to appear as an encryption of zero regardless of the values of (α, β) . Similarly, setting $(a, b) = (g^v \alpha^c \bmod p, K^{(v-c) \bmod q} \beta^c \bmod p)$ will satisfy the required checks for (α, β) to appear as an encryption of one regardless of the values of (α, β) . This quirk is what enables the Cramer-Damgård-Schoenmakers technique to prove a disjunction of two predicates.

Sketch of NIZK Proof that (α, β) is an encryption of zero or one. The prover starts with whichever statement is *not* true, generates a random challenge for it, and then derives the commitment from it to be able to produce a false proof, as described in the previous paragraph. For example, if (α, β) is an encryption of one, the prover chooses a challenge c_0 and a response at random, then derives values for (a_0, b_0) so that the ‘response’ satisfies the ‘challenge’; if (α, β) is an encryption of zero, the prover chooses c_1 and a response at random and derives satisfying values for (a_1, b_1) . The prover then works on the true part. It generates a truthful commitment (a_0, b_0) like the first part of a proof of an encryption of zero (or a truthful commitment (a_1, b_1) like the first part of a proof of an encryption of one, whichever is true). A single challenge value c is

selected by hashing all commitments and baseline parameters. The prover must produce challenge values c_0 and c_1 s.t. $c = c_0 + c_1 \bmod q$. It answers its false claim with the challenge it chose in the beginning (c_0 if (α, β) is an encryption of one, c_1 if it is an encryption of zero), then derives the other challenge by subtraction ($c_1 = c - c_0 \bmod q$ or $c_0 = c - c_1 \bmod q$ as needed). It then finishes the proof of the true part as usual, using this challenge. An observer can see that one of the claims must be true but cannot tell which.

Sketch of NIZK Proof that (α, β) is an encryption of an integer ℓ such that $0 \leq \ell \leq R$. The above method to prove that (α, β) is an encryption of zero or one can be extended to more than two values in an analogous way. Now, the prover needs to produce false proofs that verify for all values that are not encrypted in the ciphertext. This means that for each i with $0 \leq i \leq R$ and $i \neq \ell$, the prover selects a challenge value c_i and a response at random and uses them to create commitments (a_i, b_i) that make the response verify the challenge for the false claim that (α, β) is an encryption of the value $i \neq \ell$. For the true assertion that (α, β) is an encryption of ℓ , the prover makes honest commitments (a_ℓ, b_ℓ) . At this point, there exist commitments (a_i, b_i) for all $0 \leq i \leq R$ and a single challenge value c is selected by hashing all these commitments. The remaining challenge value c_ℓ is then uniquely determined by the constraint $c = (c_0 + \dots + c_\ell + \dots + c_R) \bmod q$ and can be computed by subtracting from c all other challenge values that were chosen for the false proofs. The challenge c_ℓ must be used to answer the true assertion that (α, β) is an encryption of ℓ . The proof now consists of all challenge and response pairs, including the honestly generated one for the true assertion and R faked pairs for the false assertions.

3.3.5 Details for proofs of ballot correctness

This subsection begins with the special but common case where a selection is either a 0 or a 1, which means that the range bound is $R = 1$. Both cases—unselected (encryption of 0) and selected (encryption of 1) are spelled out in detail before the general case range proof is specified. An ElectionGuard implementation can directly implement the general range proof and use it in the special case $R = 1$. There is no need to implement the case $R = 1$ separately.

Unselected option. To encrypt an “unselected” option on a ballot, a pseudo-random nonce $\xi \in \mathbb{Z}_q$ for the selection is derived from the ballot nonce ξ_B as described in Section 3.3.2, and an encryption of zero is formed as $(\alpha, \beta) = (g^\xi \bmod p, K^\xi \bmod p)$.

NIZK Proof: Proves that (α, β) is an encryption of zero or one.
(Requires knowledge of encryption nonce ξ for which (α, β) is an encryption of zero.)

To create the proof that (α, β) is an encryption of a zero or a one, randomly select u_0, u_1 , and c_1 from \mathbb{Z}_q and compute

$$(a_0, b_0) = (g^{u_0} \bmod p, K^{u_0} \bmod p) \tag{26}$$

and

$$(a_1, b_1) = (g^{u_1} \bmod p, K^{u_1 - c_1} \bmod p). \tag{27}$$

A challenge value c is formed by hashing the extended base hash H_E together with $\alpha, \beta, a_0, b_0, a_1$ and b_1 , namely

$$c = H(\mathbb{H}_E; \mathbf{0x21}, K, \alpha, \beta, a_0, b_0, a_1, b_1). \quad (28)$$

The proof consists of the four values c_0 , c_1 , v_0 , and v_1 which are computed as follows.

$$c_0 = (c - c_1) \bmod q, \quad (29)$$

$$v_0 = (u_0 - c_0 \cdot \xi) \bmod q, \quad (30)$$

$$v_1 = (u_1 - c_1 \cdot \xi) \bmod q. \quad (31)$$

These computed values satisfy the proof equations as follows.

$$c = ((c - c_1) + c_1) \bmod q = (c_0 + c_1) \bmod q, \quad (32)$$

$$(g^{v_0} \alpha^{c_0}, K^{v_0} \beta^{c_0}) \equiv_p (g^{u_0 - c_0 \xi} g^{c_0 \xi}, K^{u_0 - c_0 \xi} K^{c_0 \xi}) \equiv_p (a_0, b_0), \quad (33)$$

$$(g^{v_1} \alpha^{c_1}, K^{v_1 - c_1} \beta^{c_1}) \equiv_p (g^{u_1 - c_1 \xi} g^{c_1 \xi}, K^{u_1 - c_1 \xi - c_1} K^{c_1 \xi}) \equiv_p (a_1, b_1). \quad (34)$$

Selected option. To encrypt a “selected” option on a ballot, a pseudo-random nonce $\xi \in \mathbb{Z}_q$ for the selection is derived from the ballot nonce ξ_B as described in Section 3.3.2, and an encryption of one is formed as $(\alpha, \beta) = (g^\xi \bmod p, K^{\xi+1} \bmod p)$.

NIZK Proof: Proves that (α, β) is an encryption of zero or one.
(Requires knowledge of encryption nonce ξ for which (α, β) is an encryption of one.)

To create the proof that (α, β) is an encryption of a zero or a one, randomly select u_0 , u_1 , and c_0 from \mathbb{Z}_q and compute and

$$(a_0, b_0) = (g^{u_0} \bmod p, K^{u_0 + c_0} \bmod p) \quad (35)$$

and

$$(a_1, b_1) = (g^{u_1} \bmod p, K^{u_1} \bmod p). \quad (36)$$

A challenge value c is formed by hashing the extended base hash \mathbb{H}_E together with α , β , a_0 , b_0 , a_1 , and b_1 , namely

$$c = H(\mathbb{H}_E; \mathbf{0x21}, K, \alpha, \beta, a_0, b_0, a_1, b_1). \quad (37)$$

The proof consists of the four values c_0 , c_1 , v_0 , and v_1 which are computed as follows.

$$c_1 = (c - c_0) \bmod q, \quad (38)$$

$$v_0 = (u_0 - c_0 \cdot \xi) \bmod q, \quad (39)$$

$$v_1 = (u_1 - c_1 \cdot \xi) \bmod q. \quad (40)$$

These computed values satisfy the proof equations as follows.

$$c = (c_0 + (c - c_0)) \bmod q = (c_0 + c_1) \bmod q, \quad (41)$$

$$(g^{v_0} \alpha^{c_0}, K^{v_0} \beta^{c_0}) \equiv_p (g^{u_0 - c_0 \xi} g^{c_0 \xi}, K^{u_0 - c_0 \xi} (K^{\xi+1})^{c_0}) \equiv_p (g^{u_0}, K^{u_0 + c_0}) \equiv_p (a_0, b_0), \quad (42)$$

$$(g^{v_1} \alpha^{c_1}, K^{v_1 - c_1} \beta^{c_1}) \equiv_p (g^{u_1 - c_1 \xi} g^{c_1 \xi}, K^{u_1 - c_1 \xi - c_1} (K^{\xi+1})^{c_1}) \equiv_p (a_1, b_1). \quad (43)$$

The remainder of this section specifies the general case of a range proof used for encrypting a selection. It is a straightforward generalization of the method for proofs of encryptions of 0 or 1 described above. Any zero-knowledge proof that asserts a ciphertext (α, β) encrypts an integer value σ in the range $0, 1, \dots, R$ for some integer R is computed in detail as follows.

NIZK Proof: Proves that (α, β) is an encryption of an integer in the range $0, 1, \dots, R$. (Requires knowledge of encryption nonce ξ for which (α, β) is an encryption of ℓ .)

A disjunctive Chaum-Pedersen range proof of (α, β) being an encryption of an integer in the range $0, 1, \dots, R$ is produced as follows. First, for each $0 \leq j \leq R$, a random $u_j \in \mathbb{Z}_q$ is selected. The commitment

$$(a_\ell, b_\ell) = (g^{u_\ell} \bmod p, K^{u_\ell} \bmod p) \quad (44)$$

for $j = \ell$ is computed from u_ℓ alone. In addition, for each $0 \leq j \leq R$ with $j \neq \ell$, random $c_j \in \mathbb{Z}_q$ are selected and commitments are computed as

$$(a_j, b_j) = (g^{u_j} \bmod p, K^{t_j} \bmod p), \quad (45)$$

where $t_j = (u_j + (\ell - j)c_j) \bmod q$. Next, all the a_j, b_j values are hashed together with the ciphertext and the election's extended base hash H_E to form a pseudo-random challenge

$$c = H(H_E; 0x21, K, \alpha, \beta, a_0, b_0, a_1, b_1, \dots, a_R, b_R). \quad (46)$$

The remaining challenge value c_ℓ is determined as

$$c_\ell = \left(c - \sum_{j \neq \ell} c_j \right) \bmod q = (c - (c_0 + \dots + c_{\ell-1} + c_{\ell+1} + \dots + c_R)) \bmod q. \quad (47)$$

Finally, responses are computed for all $0 \leq j \leq R$ as

$$v_j = (u_j - c_j \xi) \bmod q \quad (48)$$

and the proof consists of the challenge values c_0, c_1, \dots, c_R and the response values v_0, v_1, \dots, v_R . Note that a range proof can be performed directly by the entity performing the public key encryption of a ballot without access to the decryption key(s). All that is required is the nonce ξ used for the selection encryption.

Specializing the general range proof with $R = 1$ and $\ell = 0$ is identical to the proof of the unselected option above, i.e. the ciphertext is an encryption of 0 or 1 given that it is an encryption of 0. Likewise, setting $R = 1$ and $\ell = 1$ is identical to the proof of the selected option.

Note 3.4. Because all the exponentiations required to encrypt and prove ballot components have a base of either g or K , implementations can be optimized by pre-computing tables of powers of these two bases. These tables can be further optimized by computing and storing the table values in Montgomery form.

Verification 5 (Well-formedness of selection encryptions)

For each selectable option on each cast ballot, an election verifier must compute the values

- (5.1) $a_j = g^{v_j} \cdot \alpha^{c_j} \bmod p$ for all $0 \leq j \leq R$,
- (5.2) $b_j = K^{w_j} \cdot \beta^{c_j} \bmod p$, where $w_j = (v_j - jc_j) \bmod q$ for all $0 \leq j \leq R$,
- (5.3) $c = H(\text{H}_E; \text{0x21}, K, \alpha, \beta, a_0, b_0, a_1, b_1, \dots, a_R, b_R)$,

where R is the option selection limit. An election verifier must then confirm the following:

- (5.A) The given values α and β are in the set \mathbb{Z}_p^r .
(A value x is in \mathbb{Z}_p^r if and only if x is an integer such that $0 \leq x < p$ and $x^q \bmod p = 1$.)
- (5.B) The given values c_j each satisfy $0 \leq c_j < 2^{256}$ for all $0 \leq j \leq R$.
- (5.C) The given values v_j are each in the set \mathbb{Z}_q for all $0 \leq j \leq R$.
(A value x is in \mathbb{Z}_q if and only if x is an integer such that $0 \leq x < q$.)
- (5.D) The equation $c = (c_0 + c_1 + \dots + c_R) \bmod q$ is satisfied.

3.3.6 Encryption of contest data

For each contest, ElectionGuard maintains a contest data field. The information held in this field must be encoded in a byte array D of a fixed length of $32 \cdot b_D$ bytes, where b_D is the smallest value large enough for D to capture all additional information from the ballot that needs to be retained and stored in encrypted form as part of the encrypted ballot. This information consists of any text written into one or more write-in text fields, information about overvotes, undervotes, and null votes, and possibly other data about voter selections. This single data field containing all such data associated with the contest is encrypted with the election public key in a single hashed ElGamal encryption as follows.

Write the contest data field as a concatenation of blocks

$$D = D_1 \parallel D_2 \parallel \dots \parallel D_{b_D}, \quad (49)$$

where the D_i , $1 \leq i \leq b_D$, consist of 32 bytes each. To encrypt it with the public election key K , a pseudo-random nonce ξ is derived from the ballot nonce ξ_B and the contest label Λ as

$$\xi = H(\text{H}_E; \text{0x20}, \xi_B, \text{ind}_c(\Lambda), \text{"contest_data"}) \quad (50)$$

to compute $(\alpha, \beta) = (g^\xi \bmod p, K^\xi \bmod p)$. A 256-bit secret key is derived as the hash

$$k = H(\text{H}_E; \text{0x22}, K, \alpha, \beta). \quad (51)$$

Next, a KDF in counter mode³³ based on HMAC is used to generate a MAC key k_0 and encryption keys $k_1 \parallel k_2 \parallel \dots \parallel k_{b_D}$ by computing

$$k_i = \text{HMAC}(k, \text{b}(i, 4) \parallel \text{Label} \parallel \text{0x00} \parallel \text{Context} \parallel \text{b}((b_D + 1) \cdot 256, 4)), \quad (52)$$

³³NIST (2022) *Recommendation for Key Derivation Using Pseudorandom Functions*. In: SP 800-108r1 <https://csrc.nist.gov/publications/detail/sp/800-108/rev-1/final>. Note that the secret session key k changes for every encryption because a fresh encryption nonce ξ is pseudo-randomly generated every time. The second input to HMAC consists of the byte encoding of the counter i , the UTF-8 encoding of the string “data_enc_keys” as the fixed Label field specified in SP 800-108r1, the separating 00 byte, the UTF-8 encoding of the string “contest_data” as the Context field, and the 4-byte encoding of the length of the total key material output in bits.

for $0 \leq i \leq b_D$. The label and context byte arrays are $\text{Label} = \mathbf{b}(\text{"data_enc_keys"}, 13)$ and $\text{Context} = \mathbf{b}(\text{"contest_data"}, 12) \parallel \mathbf{b}(\text{ind}_c(\Lambda), 4)$. Each k_i is a 256-bit key, and $\mathbf{b}(i, 4)$ and $\mathbf{b}((b_D + 1) \cdot 256, 4)$ are byte arrays of the fixed length of 4 bytes that encode the integers i and $(b_D + 1) \cdot 256$. Therefore, i and $(b_D + 1) \cdot 256$ must be less than 2^{32} , i.e. $0 \leq i < 2^{32}$ and $1 \leq b_D + 1 < 2^{24}$.

The ciphertext encrypting D is $C_D = (C_0, C_1, C_2)$, where

$$C_0 = \alpha = g^\xi \bmod p, \quad (53)$$

$$C_1 = D_1 \oplus k_1 \parallel D_2 \oplus k_2 \parallel \cdots \parallel D_{b_D} \oplus k_{b_D}, \quad (54)$$

$$C_2 = \text{HMAC}(k_0, C_0 \parallel C_1). \quad (55)$$

The component C_1 is computed by bitwise XOR (here denoted by \oplus) of each data block D_i with the corresponding key k_i . Component C_2 is a message authentication code.

3.3.7 Contest data

The contest data can contain different kinds of information such as undervote, null vote, and overvote information together with the corresponding selections, the text captured for write-in options and other data associated to the contest.

Overvote, undervote, and null vote information. The first part of the contest data field for a contest contains information on overvotes, undervotes, or null votes.

If the number of selections in a contest exceeds the selection limit, an overvote has occurred and the votes in this contest are invalid. To not affect the tallies, ElectionGuard handles an overvote situation as follows.

1. All selectable options in the contest are set to zero, i.e. ElectionGuard generates encryptions of zero (unselected option) for all selectable options.
2. An indicator that an overvote has occurred and the list of selected options are added to the contest data field.

An undervote occurs if the number of selections is less than the selection limit. A null vote occurs if no selection has been made. In both cases, an indicator for an undervote or a null vote is added to the contest data field, respectively.

Contest write-in data. For all write-in selections in a contest, ElectionGuard captures the text written into each write-in field. This text is collected in the contest data field for that contest, after any information about overvotes, undervotes or null votes.

Padding to fixed length. The collected contest data is encoded as a byte array of length $32 \cdot b_D$ bytes. If the encoded data is shorter, it is padded to the fixed length of $32 \cdot b_D$ bytes with an array of 00 bytes of the appropriate length at the end. If it is longer it must be truncated to have exactly $32 \cdot b_D$ bytes.

3.3.8 Proof of satisfying the contest selection limit

The final step in proving that a ballot is well-formed is demonstrating that the selection limits for each contest have not been exceeded. This is accomplished by homomorphically combining all of the (α_i, β_i) values for a contest by forming the aggregate contest encryption $(\bar{\alpha}, \bar{\beta}) = (\prod_i \alpha_i \bmod p, \prod_i \beta_i \bmod p)$ and proving that $(\bar{\alpha}, \bar{\beta})$ is an encryption of a value that is not larger than the total number of votes allowed for that contest (usually one). The simplest way to complete this proof is to combine all of the pseudo-random nonces ξ_i that were used to form each $(\alpha_i, \beta_i) = (g^{\xi_i} \bmod p, K^{\xi_i + \sigma_i} \bmod p)$. The aggregate nonce $\xi = \sum_i \xi_i \bmod q$ matches the aggregate contest encryption as $(\bar{\alpha}, \bar{\beta}) = (\prod_i \alpha_i \bmod p, \prod_i \beta_i \bmod p) = (g^\xi \bmod p, K^{\xi + \ell} \bmod p)$ – where $0 \leq \ell \leq L$ and L is the selection limit for the contest.

NIzk Proof: Proves that $(\bar{\alpha}, \bar{\beta})$ is an encryption of an integer in the range $0, 1, \dots, L$. (Requires knowledge of aggregate encryption nonce ξ for which $(\bar{\alpha}, \bar{\beta})$ is an encryption of ℓ .)

This proof is a disjunctive Chaum-Pedersen range proof as described in detail in Section 3.3.5. The range bound L is equal to the contest selection limit and the ciphertext (α, β) from Section 3.3.5 is replaced by the aggregate contest encryption $(\bar{\alpha}, \bar{\beta})$.

Commitments (a_i, b_i) and challenge values c_j for $j \neq \ell$ are computed as shown in Section 3.3.5. The challenge value then is computed as

$$c = H(\text{H}_E; 0x21, K, \bar{\alpha}, \bar{\beta}, a_0, b_0, a_1, b_1, \dots, a_L, b_L). \quad (56)$$

The remaining challenge value c_ℓ can then be determined and the responses v_i computed as shown in Section 3.3.5. The proof consists of the challenge values c_0, c_1, \dots, c_L and the response values v_0, v_1, \dots, v_L . Note that all of the above proofs can be performed directly by the entity performing the public key encryption of a ballot without access to the decryption key(s). All that is required is the aggregate nonce ξ , i.e. the sum of the nonces ξ_i used for the individual selection encryptions.

Verification 6 (Adherence to vote limits)

For each contest on each cast ballot, an election verifier must compute the contest totals

$$(6.1) \quad \bar{\alpha} = \prod_i \alpha_i \bmod p,$$

$$(6.2) \quad \bar{\beta} = \prod_i \beta_i \bmod p,$$

where the (α_i, β_i) represent all possible selections for the contest, as well as the values

$$(6.3) \quad a_j = g^{v_j} \cdot \bar{\alpha}^{c_j} \bmod p \text{ for all } 0 \leq j \leq L,$$

$$(6.4) \quad b_j = K^{w_j} \cdot \bar{\beta}^{c_j} \bmod p, \text{ where } w_j = (v_j - jc_j) \bmod q \text{ for all } 0 \leq j \leq L,$$

$$(6.5) \quad c = H(\text{H}_E; 0x21, K, \bar{\alpha}, \bar{\beta}, a_0, b_0, a_1, b_1, \dots, a_L, b_L),$$

where L is the contest selection limit. An election verifier must then confirm the following:

(6.A) The given values α_i and β_i are each in \mathbb{Z}_p^* .

(6.B) The given values c_j each satisfy $0 \leq c_j < 2^{256}$.

(6.C) The given values v_j are each in \mathbb{Z}_q for all $0 \leq j \leq L$.

(6.D) The equation $c = (c_0 + c_1 + \dots + c_L) \bmod q$ is satisfied.

3.4 Confirmation codes

Upon completion of the encryption of each ballot, a *confirmation code* is prepared for each voter. The code is a hash value (an output of the function H as specified in detail in Section 5) whose inputs must include the encrypted ballot and the extended base hash code H_E . Inputs to the hash may optionally include other information such as an identifier for the voting device, the location of the voting device, as well as the date and time that the ballot was encrypted.

The data that constitutes the encrypted ballot for the purpose of computing its confirmation code consists of all encrypted selections on that ballot. I.e. it includes a ciphertext (α, β) for each selection in each contest on the ballot. In detail, a confirmation code for a given ballot B is computed as follows. The number of contests on the ballot B is denoted by m_B and contests have a specified order defined in the election manifest file. Therefore, each contest has a unique sequence order l , where $1 \leq l \leq m_B$.

3.4.1 Contest hash

For each contest, ElectionGuard computes a *contest hash*, which is a hash value of an input that contains all selection option encryptions in the contest. The encryptions are hashed in order specified by the election manifest file. If $E_i = (\alpha_i, \beta_i)$ denotes the ciphertext encrypting the voter's choice for the i -th option ($1 \leq i \leq m$) of the l -th contest ($1 \leq l \leq m_B$), the contest hash value χ_l is computed from the contest label Λ_l , the election public key K , and the sequence of ciphertexts E_1, E_2, \dots, E_m as

$$\chi_l = H(H_E; 0x23, \text{ind}_c(\Lambda_l), K, \alpha_1, \beta_1, \alpha_2, \beta_2, \dots, \alpha_m, \beta_m). \quad (57)$$

3.4.2 Confirmation code

The *ballot confirmation code* is a *ballot hash* $H(B)$ that is computed from all contest hashes as

$$H(B) = H(H_E; 0x24, \chi_1, \chi_2, \dots, \chi_{m_B}, B_{\text{aux}}). \quad (58)$$

The input contains the contest hashes in the order of the contests as specified in the election manifest file. Besides the contest hashes, there is an additional, auxiliary input byte array B_{aux} . Its composition is specified in the election manifest file depending on the desired properties for the confirmation codes. The byte array B_{aux} can be used to include strings that encode the above mentioned optional information on the voting device such as a device identifier and its location, or the date and time that the ballot was encrypted. It can also contain the confirmation code of a previous ballot to enable ballot chaining as described below. If ballot chaining is not used, the manifest file may specify that this additional input is not used, in which case B_{aux} can be omitted (or treated as the empty byte array).

3.4.3 Ballot chaining

For a fixed voting device, the additional input to the confirmation code $H_j = H(B_j)$ of the j -th ballot voted on this device can also include the confirmation code $H_{j-1} = H(B_{j-1})$ of the previous

ballot B_{j-1} . When this is included, confirmation codes are part of a hash chain, which is initialized with

$$H_0 = H(H_E; 0x24, B_{aux,0}), \quad (59)$$

where $B_{aux,0}$ must contain at least a unique voting device identifier and possibly other voting device information as described above and as specified in the election manifest file.

The confirmation code for the j -th ballot when $j > 0$ is then computed as above via $H(B_j) = H(H_E; 0x24, \chi_1, \chi_2, \dots, \chi_{m_B}, B_{aux,j})$ and the additional byte array

$$B_{aux,j} = H_{j-1} \parallel B_{aux,0} \quad (60)$$

now must contain the confirmation code H_{j-1} of the previous ballot B_{j-1} (or the initialization code H_0 when computing $H_1 = H(B_1)$) and the voting device information by including $B_{aux,0}$. If $B_{aux,j}$ is allowed to contain more than one prior confirmation code, a tree of hash dependencies can be formed. The auxiliary byte array $B_{aux,j}$ should be stored in the encrypted ballot.

The chain should be closed at the end of an election by forming and publishing

$$\bar{H} = H(H_E; 0x24, \bar{B}_{aux}), \quad (61)$$

using $\bar{B}_{aux} = H(B_\ell) \parallel B_{aux,0} \parallel \mathbf{b}(\text{"CLOSE"}, 5)$, where $H(B_\ell)$ is the final confirmation code in the chain. The benefit of a chain is that it makes it more difficult for a malicious insider to selectively delete ballots and confirmation codes after an election without detection.

Verification 7 (Validation of confirmation codes)

An election verifier must confirm the following for each ballot B .

- (7.A) The contest hash χ_l for the contest with contest index l for all $1 \leq l \leq m_B$ has been correctly computed from the contest selection encryptions (α_i, β_i) as

$$\chi_l = H(H_E; 0x23, l, K, \alpha_1, \beta_1, \alpha_2, \beta_2 \dots, \alpha_m, \beta_m).$$

- (7.B) The ballot confirmation code $H(B)$ has been correctly computed from the contest hashes and if specified in the election manifest file from the additional byte array B_{aux} as

$$H(B) = H(H_E; 0x24, \chi_1, \chi_2, \dots, \chi_{m_B}, B_{aux}).$$

An election verifier must also verify the following.

- (7.C) There are no duplicate confirmation codes, i.e. among the set of submitted (cast and challenged) ballots, no two have the same confirmation code.

Additionally, if the election manifest file specifies a hash chain, an election verifier must confirm the following for each voting device.

- (7.D) The initial hash code H_0 satisfies $H_0 = H(H_E; 0x24, B_{aux,0})$ and $B_{aux,0}$ contains the unique voting device information.
- (7.E) For all $1 \leq j \leq \ell$, the additional input byte array used to compute $H_j = H(B_j)$ is equal to $B_{aux,j} = H(B_{j-1}) \parallel B_{aux,0}$.
- (7.F) The final additional input byte array is equal to $\bar{B}_{aux} = H(B_\ell) \parallel B_{aux,0} \parallel \mathbf{b}(\text{"CLOSE"}, 5)$ and $H(B_\ell)$ is the final confirmation code on this device.
- (7.G) The closing hash is correctly computed as $\bar{H} = H(H_E; 0x24, \bar{B}_{aux})$.

Ballot casting or challenging. Once in possession of a confirmation code (*and never before*), a voter is afforded an option to either cast the associated ballot or challenge it and restart the ballot preparation process. The precise mechanism for voters to make these selections may vary depending upon the instantiation, but this choice would ordinarily be made immediately after a voter is presented with the confirmation code, and the status of the ballot would be undetermined until the decision is made. It is possible, for instance, for a voter to make the decision directly on the voting device, or a voter may instead be afforded an option to deposit the ballot in a receptacle or to take it to a poll worker to be challenged. For vote-by-mail scenarios, a voter can be sent (hashes of) two complete sets of encryptions for each selectable option and can effect a ballot challenge implicitly by choosing which encryptions to return.

3.5 Ballot Aggregation

At the conclusion of voting, all of the ballot encryptions are published in the election record together with the proofs that the ballots are well-formed. Additionally, all of the encryptions of each option are homomorphically combined to form an encryption of the sum of the values that were individually encrypted. The encryptions (α_i, β_i) of each individual option are combined by forming the product

$$(A, B) = \left(\prod_i \alpha_i \bmod p, \prod_i \beta_i \bmod p \right). \quad (62)$$

This aggregate encryption (A, B) , which represents an encryption of the tally of that option, is published in the election record for each option.

Verification 8 (Correctness of ballot aggregation)

An election verifier must confirm for each option in each contest in the election manifest that the aggregate encryption (A, B) satisfies

$$(8.A) \quad A = \left(\prod_j \alpha_j \right) \bmod p,$$

$$(8.B) \quad B = \left(\prod_j \beta_j \right) \bmod p,$$

where the (α_j, β_j) are the corresponding encryptions on all cast ballots in the election record.

3.6 Verifiable Decryption

To decrypt an aggregate encryption (A, B) (or an individual encryption such as one on a challenged ballot), guardians work together in a protocol, where each guardian produces a partial decryption and contributes to an accumulated Chaum-Pederson proof of correct decryption.

As long as at least k guardians are present for decryption, the partial decryptions produced by the guardians are used to compute the value

$$M = A^s \bmod p, \quad (63)$$

without ever computing the secret key s . This value is then used to obtain

$$T = B \cdot M^{-1} \bmod p. \quad (64)$$

This T has the property that $T = K^t \bmod p$ where t is the tally of the associated option.

In general, computation of this tally value t is computationally intractable. However, in this application, t is relatively small—usually bounded by the number of votes cast. This tally value t can be determined from T by exhaustive search, by precomputing a table of all possible T values in the allowable range and then performing a single look-up, or by a combination in which some exponentiations are precomputed and a small search is used to find the value of t (e.g., a partial table consisting of $K^{100} \bmod p$, $K^{200} \bmod p$, $K^{300} \bmod p$, ... is precomputed and the value T is repeatedly divided (or multiplied) by K until a value is found that is in the partial table). The value t is published in the election record, and verifiers should check both that $T = K^t \bmod p$ and that $B = T \cdot M \bmod p$.

3.6.1 Partial decryption by available guardians

During the key generation process (as described in Section 3.2.2), each guardian G_i receives from each other guardian G_j a share $P_j(i)$ of G_j 's secret key s_j ($1 \leq i, j \leq n$ and $i \neq j$). Guardian G_i also computes its own share $P_i(i)$ and is therefore in possession of all shares $P_j(i)$ for $1 \leq j \leq n$. Either at that time or at a later time but before G_i participates in any joint decryption process, G_i computes

$$P(i) = \left(\sum_{j=1}^n P_j(i) \right) \bmod q = (P_1(i) + P_2(i) + \cdots + P_n(i)) \bmod q. \quad (65)$$

The value $P(i)$ is G_i 's share of the implicit secret key $s = (s_1 + s_2 + \cdots + s_n) \bmod q$.³⁴

Each guardian G_i available for decryption uses the share $P(i)$ to compute the partial decryption

$$M_i = A^{P(i)} \bmod p \quad (66)$$

that contributes to computing the value $M = A^s \bmod p$.

3.6.2 Combination of partial decryptions.

The partial decryptions are combined to obtain M with the help of the Lagrange coefficients that correspond to the available guardians. Let $U \subseteq \{1, 2, \dots, n\}$ be the set of indices such that $\{G_i : i \in U\}$ is the set of available guardians. For decryption to succeed, a quorum of at least k guardians is needed, i.e. $|U| = h \geq k$. All available guardians participate in the reconstruction even if $h > k$, i.e., there are more guardians available than strictly necessary for having a quorum.

The Lagrange coefficients corresponding to U are computed as

$$w_i = \left(\prod_{\ell \in (U \setminus \{i\})} \frac{\ell}{\ell - i} \right) \bmod q. \quad (67)$$

³⁴Guardian G_i 's secret key s_i is the constant coefficient of the polynomial $P_i(x)$ and therefore the implicit secret key $s = (s_1 + s_2 + \cdots + s_n) \bmod q$ is the constant coefficient of the polynomial $P(x) = (P_1(x) + P_2(x) + \cdots + P_n(x)) \bmod q$. The collection of the $P(i)$ are the shares for a k -out-of- n sharing of the secret key s . It is important to note that this secret key s is never actually computed; instead each guardian uses its share of the implicit secret key s to form a share of any decryption that needs to be performed.

These coefficients are used to combine the M_i for $i \in U$ provided by the available guardians to obtain

$$M = \prod_{i \in U} (M_i)^{w_i} \bmod p. \quad (68)$$

Note 3.5. The decryption $M = A^s \bmod p$ can be computed as shown in Equation (68) because $s = (\sum_{i \in U} w_i P(i)) \bmod q$. Likewise, a missing secret s_j could be computed directly as $s_j = (\sum_{i \in U} w_i P_j(i)) \bmod q$. However, it is preferable to not release the secret s (or any of the missing secrets) and instead only release the partial decryptions that the secret would have produced. This prevents the secret from being used for additional decryptions without the cooperation of at least k guardians.

As an example, consider an election with five guardians and a threshold of three. If two guardians are missing at the time of decryption, the remaining three can perform any required decryptions as described in the text above. If, instead, they take the shortcut of simply reconstructing and then using the two missing secrets, then any of the three could, at a later time, use these missing secrets together with its own secret to perform additional decryptions without the cooperation of any other guardian.

3.6.3 Proof of correctness

The available guardians work together to produce a Chaum-Pedersen proof that M was computed correctly, which means that $M = A^s \bmod p$. The proof is computed as follows.

NIZK Proof: The available guardians jointly prove that they have shared knowledge of $s \in \mathbb{Z}_q$ such that $M = A^s \bmod p$ and $K = g^s \bmod p$.

Each available guardian G_i , $i \in U$, selects a random value u_i in \mathbb{Z}_q and commits to the pair

$$(a_i, b_i) = (g^{u_i} \bmod p, A^{u_i} \bmod p). \quad (69)$$

The a_i and b_i obtained from each guardian are used to compute the accumulated commitments as

$$a = \left(\prod_{i \in U} a_i \right) \bmod p, \quad b = \left(\prod_{i \in U} b_i \right) \bmod p. \quad (70)$$

This means $a = g^u \bmod p$ and $b = A^u \bmod p$, where $u = (\sum_{i \in U} u_i) \bmod q$ and u is not computed explicitly.

The ciphertext (A, B) , the commitments (a, b) , and the combined value M are then hashed together with the extended base hash value H_E to form a challenge

$$c = H(H_E; 0x30, K, A, B, a, b, M). \quad (71)$$

The challenge c is adjusted by the i -th Lagrange coefficient to produce a challenge c_i for available guardian G_i ($i \in U$) as

$$c_i = (c \cdot w_i) \bmod q. \quad (72)$$

Next, each available guardian G_i responds to the challenge c_i with

$$v_i = (u_i - c_i P(i)) \bmod q. \quad (73)$$

All individual responses v_i are verified by recomputing the commitments from the responses, the individual challenge values c_i , the guardians' commitments $K_{j,m}$ to the coefficients $a_{j,m}$ of their secret sharing polynomials P_j (see 3.2.2), the ciphertext value A and the partial decryptions M_i as

$$a'_i = g^{v_i} \left(\prod_{j=1}^n \prod_{m=0}^{k-1} K_{j,m}^{i^m} \right)^{c_i} \bmod p, \quad (74)$$

$$b'_i = A^{v_i} M_i^{c_i} \bmod p \quad (75)$$

and checking that $a'_i = a_i$ and $b'_i = b_i$. If these equations are all satisfied, an accumulated response

$$v = \left(\sum_{i \in U} v_i \right) \bmod q \quad (76)$$

is computed and the decrypted value $T = B \cdot M^{-1} \bmod p$ is published along with the proof (c, v) in the election record. Note that $v = (u - c \cdot s) \bmod q$.

Verification 9 (Correctness of tally decryptions)

For each option in each contest on each tally, an election verifier must compute the values

$$(9.1) \quad M = B \cdot T^{-1} \bmod p,$$

$$(9.2) \quad a = g^v \cdot K^c \bmod p,$$

$$(9.3) \quad b = A^v \cdot M^c \bmod p.$$

An election verifier must then confirm the following:

(9.A) The given value v is in the set \mathbb{Z}_q .

(9.B) The challenge value c satisfies $c = H(H_E; 0x30, K, A, B, a, b, M)$.

Tally verification. The final step is to verify the tallies themselves.

Verification 10 (Validation of correct decryption of tallies)

An election verifier must confirm the following equations for each option in each contest in the election manifest.

(10.A) $T = K^t \pmod p$.

An election verifier must also confirm that the text labels listed in the election record tallies match the corresponding text labels in the election manifest. For each contest in a decrypted tally, an election verifier must confirm the following.

(10.B) The contest text label occurs as a contest label in the list of contests in the election manifest.

(10.C) For each option in the contest, the option text label occurs as an option label for the contest in the election manifest.

(10.D) For each option text label listed for this contest in the election manifest, the option label occurs for a option in the decrypted tally contest.

An election verifier must also confirm the following.

(10.E) For each contest text label that occurs in at least one submitted ballot, that contest text label occurs in the list of contests in the corresponding tally.

3.6.4 Decryption of contest data

For each contest, an encrypted ballot contains a ciphertext $C_E = (C_0, C_1, C_2)$ encrypting contest data such as overvote, undervote, and null vote information and write-in text fields. The ciphertext has been generated as described in Section 3.3.6. Such data may need to be decrypted if the tallies record a significant number of votes in write-in selections. Also, all contest data fields on each challenged ballot are decrypted.

Decryption can be done by a quorum of guardians similarly to the decryption of tallies explained above. Each available guardian G_i , $i \in U$, computes a partial decryption

$$m_i = C_0^{P(i)} \pmod p \tag{77}$$

using its precomputed share $P(i) = \sum_{j=1}^n P_j(i) \pmod q$. The partial decryptions are combined using the Lagrange coefficients for the set U of available guardians to obtain

$$\beta = \left(\prod_{i \in U} m_i^{w_i} \right) \pmod p \tag{78}$$

Again, the available guardians work together to produce and publish the following proof.

NIZK Proof: The available guardians jointly prove that they have shared knowledge of $s \in \mathbb{Z}_q$ such that $\beta = C_0^s \pmod p$ and $K = g^s \pmod p$.

This proof is exactly the same as the one in Section 3.6.3, where the ciphertext (A, B) is replaced by the contest data ciphertext (C_0, C_1, C_2) as follows. Guardian G_i selects a random value u_i in \mathbb{Z}_q and commits to the pair

$$(a_i, b_i) = (g^{u_i} \pmod p, C_0^{u_i} \pmod p). \tag{79}$$

The values a_i and b_i are accumulated into

$$a = \left(\prod_{i \in U} a_i \right) \bmod p, \quad b = \left(\prod_{i \in U} b_i \right) \bmod p \quad (80)$$

and the joint challenge value c is obtained as

$$c = H(\text{H}_E; \text{0x31}, K, C_0, C_1, C_2, a, b, \beta). \quad (81)$$

Each available guardian G_i responds to its challenge $c_i = (c \cdot w_i) \bmod q$ with

$$v_i = (u_i - c_i P(i)) \bmod q. \quad (82)$$

As above, all proof parts v_i are verified by recomputing commitments

$$a'_i = g^{v_i} \left(\prod_{j=1}^n \prod_{m=0}^{k-1} K_{j,m}^{i^m} \right)^{c_i} \bmod p, \quad (83)$$

$$b'_i = C_0^{v_i} m_i^{c_i} \bmod p. \quad (84)$$

If $a'_i = a_i$ and $b'_i = b_i$, an accumulated response

$$v = \left(\sum_{i \in U} v_i \right) \bmod q \quad (85)$$

is computed and the decryption value β is published along with the proof (c, v) .

Then decryption proceeds by computing the key $k = H(\text{H}_E; \text{0x22}, K, C_0, \beta)$ and the MAC key

$$k_0 = \text{HMAC}(k, \text{b}(0, 4) \parallel \text{Label} \parallel \text{0x00} \parallel \text{Context} \parallel \text{b}((b_D + 1) \cdot 256, 4)) \quad (86)$$

with $\text{Label} = \text{b}(\text{"data_enc_keys"}, 13)$ and $\text{Context} = \text{b}(\text{"contest_data"}, 12) \parallel \text{b}(\text{ind}_c(\Lambda), 4)$. When it is the case that $C_2 = \text{HMAC}(k_0, C_0 \parallel C_1)$, the encryption keys k_1, k_2, \dots, k_{b_D} are computed as

$$k_i = \text{HMAC}(k, \text{b}(i, 4) \parallel \text{Label} \parallel \text{0x00} \parallel \text{Context} \parallel \text{b}((b_D + 1) \cdot 256, 4)) \quad (87)$$

and the byte array D representing the contest data string is decrypted by parsing C_1 in 32-byte blocks as

$$C_1 = C_{1,1} \parallel C_{1,2} \parallel \dots \parallel C_{1,b_D} \quad (88)$$

and obtaining

$$D = C_{1,1} \oplus k_1 \parallel C_{1,2} \oplus k_2 \parallel \dots \parallel C_{1,b_D} \oplus k_{b_D}. \quad (89)$$

The byte array D can now be parsed to reveal the captured overvote, undervote, and null vote data and write-in text fields.

Verification 11 (Correctness of decryptions of contest data)

An election verifier must confirm the correct decryption of the contest data field for each contest by verifying the conditions analogous to Verification 9 for the corresponding NIZK proof with (A, B) replaced by (C_0, C_1, C_2) and M_i by m_i as follows. An election verifier must compute the following values.

$$(11.1) \quad a = g^v \cdot K^c \pmod p,$$

$$(11.2) \quad b = C_0^v \cdot \beta^c \pmod p.$$

An election verifier must then confirm the following.

(11.A) The given value v is in the set \mathbb{Z}_q .

(11.B) The challenge value c satisfies $c = H(H_E; 0x31, K, C_0, C_1, C_2, a, b, \beta)$.

3.6.5 Decryption of challenged ballots

Each and every challenged ballot must be verifiably decrypted. ElectionGuard supports two methods of decrypting challenged ballots. The first is for the guardians to use their secret keys in precisely the way they do to decrypt election tallies. The second is to simply publish the primary nonce used to encrypt each challenged ballot. Whether these nonces are available depends on how ElectionGuard is used.

Standard decryption with guardian keys. Every ballot challenged in an election is individually verifiably decrypted in exactly the same way that the aggregate ballot of tallies is decrypted. Each individual selection encryption (α, β) in each contest on the challenged ballot is jointly decrypted by the available guardians as shown at the beginning of Section 3.6 with (A, B) replaced by (α, β) . This includes the generation of the appropriate NIZK proofs to prove correctness of these decryptions M . The corresponding encrypted selection is then decrypted by combining the partial decryptions as shown above for tallies by computing

$$S = \beta \cdot M^{-1} \pmod p. \tag{90}$$

This S encrypts a vote σ as $S = K^\sigma \pmod p$ from which σ can be determined as described at the beginning of Section 3.6.³⁵

In addition, all contest data fields on a challenged ballot are decrypted as shown in the previous Subsection 3.6.4.

An election verifier must confirm the correct decryption of each challenged ballot using the same process that was used to confirm the election tallies. In particular, this means that Verification 9 must be confirmed. In order to assign unique verification steps for challenged ballots, the relevant verification steps for tallies are adjusted and listed here again.

In addition to all selection and contest data decryptions, election verifiers should confirm that a ballot is well-formed. Overall, casual observers should be able to simply view the decryptions and confirm that they match their expectations.

³⁵When, as is usually the case, σ is known to be in the set $\{0, 1\}$, the decryption can be completed simply by checking whether $S = 1$ or $S = K$.

Verification 12 (Correctness of decryptions for challenged ballots)

For each challenged ballot, for each option in each contest on the challenged ballot, an election verifier must compute the values

- (12.1) $M = \beta \cdot S^{-1} \bmod p$,
- (12.2) $a = g^v \cdot K^c \bmod p$,
- (12.3) $b = \alpha^v \cdot M^c \bmod p$.

An election verifier must then confirm the following.

- (12.A) The given value v is in the set \mathbb{Z}_q .
- (12.B) The challenge value c satisfies $c = H(H_E; 0x30, K, \alpha, \beta, a, b, M)$.

Verification 13 (Validation of correct decryption of challenged ballots)

An election verifier must confirm the correct decryption of each selection σ in each contest.

- (13.A) $S = K^\sigma \bmod p$.

An election verifier must also confirm that the challenged ballot is well-formed, i.e. for each contest on the challenged ballot, it must confirm the following.

- (13.B) For each option in the contest, the selection σ is a valid value — usually either a 0 or a 1.
- (13.C) The sum of all selections in the contest is at most the selection limit L for that contest.

An election verifier must also confirm that for each decrypted challenged ballot, the selections listed in text match the corresponding text in the election manifest.

- (13.D) The contest text label occurs as a contest label in the list of contests in the election manifest.
- (13.E) For each option in the contest, the option text label occurs as an option label for the contest in the election manifest.
- (13.F) For each option text label listed for this contest in the election manifest, the option label occurs for a option in the decrypted challenged ballot.

Finally, the following verification steps are carried out for each challenged ballot to confirm the correct decryption of contest data.

Verification 14 (Correctness of contest data decryptions for challenged ballots)

An election verifier must confirm the correct decryption of the contest data field for each contest by verifying the conditions analogous to Verification 9 for the corresponding NIZK proof with (A, B) replaced by (C_0, C_1, C_2) and M by β as follows. This means it must compute the values

- (14.1) $a = g^v \cdot K^c \bmod p$,
- (14.2) $b = C_0^v \cdot \beta^c \bmod p$.

It must then confirm the following.

- (14.A) The given value v is in the set \mathbb{Z}_q .
- (14.B) The challenge value c satisfies $c = H(H_E; 0x31, K, C_0, C_1, C_2, a, b, \beta)$.

Verifying decryption with nonces. When the primary nonce ξ_B that was used to encrypt a ballot B is available, a verifiable decryption can be produced by simply publishing this nonce. This

can drastically reduce the amount of computation required of both guardians and verifiers.

There are several scenarios in which this nonce may be available. For instance, if a ballot is challenged immediately after encryption, the device that performed the encryption may be able to provide the nonce as soon as the challenge is issued. Alternatively, the device that encrypted the ballot might encrypt the ballot's primary nonce with the election key or with another key and provide this with the encrypted ballot.

Publishing a ballot's primary nonce in an election record enables a decryption to be confirmed by simply repeating the encryption and checking for a match. The same steps could also be taken by a voter equipped with a suitable application immediately upon receiving this nonce to verify a challenge ballot without having to wait for the election record to be published.³⁶

Specifically, with possession of a primary ballot nonce ξ_B and the selections made on the ballot, a verification application can repeat the vote encryption process as described in §3.3 as follows.

- It derives the selection encryption nonce for each contest and each selection in the contest. The nonce for the i -th contest, which has label Λ_i , and the j -th selection in the contest, which has label λ_j is $\xi_{i,j}$ and is derived by a hash computation as shown in Equation (25) in Section 3.3.2.
- Using the encryption nonces $\xi_{i,j}$, it recomputes the selection encryptions as shown in Equation (24) of Section 3.3.1.
- Finally, it recomputes the confirmation code $H(B)$ as described in Section 3.4 via its contest hashes.

If the resulting confirmation code matches the confirmation code provided to the voter, then this is verification that the confirmation code corresponds to a ballot with votes for the indicated selections.

Note that confirmation codes *do not* include any of the zero-knowledge proofs of ballot correctness. So a verifier does not need to reproduce these zero-knowledge proofs. Only the actual encryptions of the selections must be regenerated and confirmed.

3.7 The Election Record

The record of an election should be a full accounting of all of the election artifacts. Specifically, it should contain the following.

- Information sufficient to uniquely identify and describe the election, such as date, location, election type, etc. (not otherwise included in the election manifest).
- The election manifest file.
- The baseline parameters:
 - primes p and q and integer r such that $p = qr + 1$ and r is not a multiple of q ,
 - a generator g of the order q multiplicative subgroup \mathbb{Z}_p^* ,
 - the number n of election guardians,
 - the quorum threshold k of guardians required to complete verification.
- The parameter base hash H_P computed from the parameters.

³⁶It is likely that a voter would only have access to a ballot's confirmation code — not its full encryption — if a verification is to be performed prior to the publication of the election record. Thus, the voter's application would need to regenerate the encrypted ballot from the primary nonce and the voter's selections and then recompute the confirmation code in order to verify its accuracy.

- The base hash value H_B computed from the above.
- The commitments from each election guardian to each of their polynomial coefficients.
- The proofs from each guardian of possession of each of the associated coefficients.
- The election public key.
- The extended base hash value H_E computed from the above.
- Every encrypted ballot prepared in the election (whether cast or challenged):
 - All of the encrypted selections on each ballot,
 - the proofs that each such value is an encryption of either zero or one,
 - the selection limit for each contest,
 - the proof that the number of selections made matches the selection limit,
 - the ballot style,
 - the device information for the device that encrypted the ballot,
 - the date and time of the ballot encryption,
 - the confirmation code produced for the ballot.
- The decryption of each challenged ballot:
 - The selections made on the ballot,
 - the plaintext representation of the selections,
 - proofs of each decryption.
- Tallies of each option in an election:
 - The encrypted tally of each option,
 - full decryptions of each encrypted tally,
 - plaintext representations of each tally,
 - proofs of correct decryption of each tally.
- Ordered lists of the ballots encrypted by each device.

An election record should be digitally signed by election administrators together with the date of the signature. The entire election record and its digital signature should be published and made available for full download by any interested individuals. Tools should also be provided for easy look up of confirmation codes by voters.

The exact organizational structure of the election record will be specified in a separate document.

4 Pre-Encrypted Ballots

In typical use, ElectionGuard ballots are encrypted after the voter selections are known. However, there are several common scenarios in which it is desirable to perform the encryption before the selections of a voter are known. These include Vote-by-Mail, pre-printed ballots for use in precincts with central count or minimal in-precinct equipment, and back-ups for precincts which ordinarily perform encryption on demand.

With pre-encrypted ballots, each possible selection on a ballot is individually encrypted in advance; and the selections made by the voter indicate which encryptions are used.

ElectionGuard requires two applications to support pre-encrypted ballots: a *ballot encrypting tool* to provide data to enable printing of blank ballots and a companion *ballot recording tool* to receive information about selections made on pre-printed ballots and produce data compatible with the ElectionGuard election record.

4.1 Format of Pre-Encrypted Ballots

Each selectable option within each contest is associated with a vector Ψ of encryptions E_j — with one encryption for each selectable option within the contest. Selectable options in a contest have a unique order determined by their option indices in increasing order. Let i be the position of the selectable option in this ordering (not necessarily identical to its option index). In its *normal* form and for a contest with m selection options, this vector

$$\Psi_{i,m} = \langle E_1, E_2, \dots, E_m \rangle \tag{91}$$

includes an encryption $E_i = (\alpha_i, \beta_i) = \text{Enc}(1; \xi_i)$ of one in the vector position i associated with the selection made and encryptions $E_j = (\alpha_j, \beta_j) = \text{Enc}(0; \xi_j)$ of zero in every other position $1 \leq j \leq m, j \neq i$, where the ξ_j are (pseudo-)random encryption nonces. For example, the vector $\Psi_{2,4}$ associated with selecting the second option in a contest with a total of $m = 4$ selection options is $\Psi_{2,4} = \langle \text{Enc}(0; \xi_1), \text{Enc}(1; \xi_2), \text{Enc}(0; \xi_3), \text{Enc}(0; \xi_4) \rangle$. This corresponds precisely to the standard ElectionGuard encryption of a vote for the same option. There is also a *null* form $\Psi_{0,m} = \langle \text{Enc}(0; \xi_1), \text{Enc}(0; \xi_2), \dots, \text{Enc}(0; \xi_m) \rangle$ which is the same form except that all values are encryptions of zero.

The principal difference between a pre-encrypted ballot and a standard ElectionGuard ballot is that while standard ElectionGuard has a single vector of encryptions for each contest, here we have a vector of encryptions for each selectable option in each contest (generally including the possibility of an undervote in which no selections are made). Another difference is that while a standard ElectionGuard contest encryption can contain multiple selections, each vector of pre-encryptions represents at most one selection. However, in a contest where a voter is allowed to make multiple selections, multiple pre-encryption vectors can be combined to form a single contest encryption vector with multiple encryptions of one that precisely matches the standard ElectionGuard format.

4.1.1 Selection hash

Each pre-encrypted vector of a pre-encrypted ballot is hashed using the ElectionGuard hash function H (specified in detail in Section 5) to form a *selection hash*. For all selectable options, i.e. for each

option at position i in the contest with $1 \leq i \leq m$, the hash value ψ_i of the selection vector $\Psi_{i,m} = \langle E_1, E_2, \dots, E_m \rangle$ is computed as

$$\psi_i = H(H_E; 0x40, K, \alpha_1, \beta_1, \alpha_2, \beta_2 \dots, \alpha_m, \beta_m), \quad (92)$$

where $E_i = (\alpha_i, \beta_i)$ is an encryption of one and $E_j = (\alpha_j, \beta_j)$ is an encryption of zero for $j \neq i$.

In a contest with a selection limit of L , an additional L null vectors are hashed to obtain

$$\psi_{m+\ell} = H(H_E; 0x40, K, \alpha_1, \beta_1, \alpha_2, \beta_2 \dots, \alpha_m, \beta_m), \quad (93)$$

where all $E_i = (\alpha_i, \beta_i)$ are encryptions of zero and $1 \leq \ell \leq L$.

4.1.2 Contest hash

All of the selection hashes within each contest will ultimately be hashed *in sorted order* to form the *contest hash* of that contest. Each contest on a ballot has a unique position determined by the position of its contest index in the list of contest indices in increasing order. The contest hash for the l -th contest (with label Λ_l) on the ballot is computed as

$$\chi_l = H(H_E; 0x41, \text{ind}_c(\Lambda_l), K, \psi_{\sigma(1)}, \psi_{\sigma(2)}, \dots, \psi_{\sigma(m+L)}), \quad (94)$$

where σ is a permutation that represents the sorting of the selection hashes. This means that contests are *not* hashed in the order given by the contest indices, but instead $\sigma(i) < \sigma(j)$ implies that $\psi_{\sigma(i)} < \psi_{\sigma(j)}$ (when hash values are interpreted as integers in big endian byte order). The sorting is required so that the order of the selection hashes $\psi_1, \psi_2, \dots, \psi_m$ does not reveal the contents of the encryptions that are used to generate the hashes.

4.1.3 Confirmation code

While contest hashes for pre-encrypted ballots are computed from selection hashes, which differs from the standard scenario for ElectionGuard described in Section 3.4, the computation of confirmation codes aligns with the previous case. The *confirmation code* $H(B)$ of a pre-encrypted ballot B is generated as the hash of all the contest hashes on the ballot in sequential order. If there are m_B contests on the ballot (in sequential order specified by their contest indices in the election manifest file), its confirmation code is computed as

$$H(B) = H(H_E; 0x42, \chi_1, \chi_2, \dots, \chi_{m_B}, B_{\text{aux}}). \quad (95)$$

Use of the auxiliary input byte array B_{aux} is the same as described in Sections 3.4.2 and 3.4.3 to enable the input of specific information of the device generating pre-encrypted ballots.

A ballot's hash will be printed directly on the ballot. Ideally, two copies of the ballot hash will be included on each ballot with one remaining permanently with the ballot and the other in an immediately adjacent location on a removable tab. These removable ballot codes are intended to be retained by voters and can be used as an identifier to allow voters to look up their (encrypted) ballots in the election record and confirm that they include the proper *short codes* as described below.

4.1.4 Short Codes

In any instantiation of pre-encrypted ballots, an additional *hash trimming function* Ω must be provided. The hash trimming function takes as its input a selection hash, and its output is a *short code*. As an example, Ω could produce the last byte of its input in a specified form. For instance, a short code representation could be a pair of hex characters, a pair of letters from a 16-letter alphabet, a letter followed by a digit, or a three-digit number. The size of a short code does not need to be a single byte, but this is a convenient choice. Different vendors or jurisdictions might choose to distinguish themselves by using their own preferred short code formats, so the details of the short code format are intentionally left open. However, Ω must be completely specified in the election manifest so that a verifier can match its functionality.

The hash trimming function Ω associates each selection on a ballot with a short code. *The short codes on a ballot need not be unique.* However, it is required that the short codes within a contest be unique. If there is a collision of short codes within a contest, the randomness should be changed to generate a new ballot. When a pre-encrypted ballot is presented to a voter, the short codes associated with each selection should be displayed beside the selection. If the ballot is cast by a voter, the short codes associated with selections made by the voter will be published as part of the election record.

Undervotes. In addition to the short codes for each possible selection, short codes are provided to indicate undervotes. (It may not be necessary to print undervote short codes on ballots.) A short code for a null vote is generated from a vector of encryptions of zero. In a contest in which the voter may select only one option there will be a single pre-encrypted null vote and associated short code to indicate that the voter did not make a selection in that contest. In general, the number of null votes and short codes for a contest should match the selection limit of the contest. So, for example, a contest in which a voter may make three selections should have three null short codes. Since the short codes corresponding to the selections made by each voter will be published in the election record, the use of null short codes allows the election record to not reveal undervotes.

4.2 The Ballot Encrypting Tool

The encrypting tool for pre-encrypted ballots takes as input parameters

- the election manifest,
- a ballot style index,
- and an optional nonce encryption key (this may be the election encryption key).

The tool produces the following outputs – which can be used to construct a single pre-encrypted ballot. (Note that it will likely be desirable to construct a wrapper that can be called to produce data for a specified number of pre-encrypted ballots.)

- An (optional) encryption of the primary nonce used to encrypt the ballot,
- a selection hash value for each possible selection in the ballot style,
- for each contest, additional null hash values corresponding in number to the contest selection limit,
- a contest hash for each contest computed from the sorted list of selection hashes and null

- hashes for that contest,
- and a confirmation code consisting of a hash of all of the contest hashes on the ballot.

The encrypting tool operates as follows.

First, it generates a 256-bit primary nonce ξ for the ballot and, if a nonce encryption key is provided, encrypts this primary nonce with the nonce encryption key provided.

Next, for each contest on the indicated ballot style, an encryption vector is produced for each selection within the contest (see Equation (91)). This encryption vector is deterministically derived from the primary nonce ξ and consists of an encryption of one in the position corresponding to the selection and an encryption of zero in all other positions in the contest (see Equation (92)). Additionally, one or more null vectors consisting entirely of encryptions of zeros are produced (again deterministically from the primary nonce) as in Equation (93). The number of null vectors should match the selection limit L of the contest. The selection hashes and null hashes (computed as described in Section 4.1.1) are then sorted numerically and hashed together (in sorted order) to produce the contest hash as shown in Equation (94) in Section 4.1.2.

Finally, the contest hashes are themselves hashed sequentially to form the ballot’s confirmation code according to Equation (95) in Section 4.1.3.

4.2.1 Deterministic nonce derivation

The process of deterministic encryption is guided by a primary nonce ξ . It is assumed that each contest has a label Λ and that within each contest each possible selection has a label λ . The nonce used within the i^{th} contest and within that the j^{th} selection vector to form the k^{th} encryption is

$$\xi_{i,j,k} = H(H_E; 0x43, \xi, \text{ind}_c(\Lambda_i), \text{ind}_o(\lambda_j), \text{ind}_o(\lambda_k)). \quad (96)$$

Note that the nonce $\xi_{i,j,k}$ will be used to encrypt a one whenever $j = k$ and a zero whenever $j \neq k$. Note also that some of the selection labels will represent null votes. If labels for null votes are not included within the manifest file, the labels “null1”, “null2”, ... should be used within each contest as needed.

The output of the encryption tool includes the (optionally encrypted) primary nonce, the selection hash corresponding to each selection on the ballot (including nulls), the contest hashes, and the confirmation code.

4.2.2 Using the Ballot Encrypting Tool

It is the responsibility of the entity that calls the encryption tool to produce short codes from selection hashes. This must be done in a deterministic, repeatable fashion using a hash trimming function Ω . As suggested above, one option is to use a human-friendly representation of the final byte (or bytes) of each selection hash. Another would be to use an ordinal integer to indicate where in the sorted order of selection hashes within each contest each selection falls. (For example, if a contest has four possible selections, the integers 1, 2, 3, 4, 5 could be placed beside each selection – including “none” to indicate the sorted position of each of the five selection hashes.) This flexibility allows vendors to distinguish themselves with different ballot presentations and for vendors and jurisdictions to choose presentations that best accommodate their voters.

Unless the short codes in a particular instantiation are quite long, it is likely that there will be occasional collisions of short codes within a contest. It is the responsibility of the entity that calls the ballot encryption tool to ensure that no ballot is printed with a short code that is repeated within a contest. If a collision is found, the caller simply discards this ballot data and calls the ballot encryption tool again. The caller may also choose to discard ballot data if a short code is repeated anywhere within a ballot or if two short codes – within a contest or across a ballot – meet some definition of similarity. The caller is free to discard data and obtain a new ballot pre-encryption as often as it likes. (Note that a malicious encryption wrapper could bias the printed ballots by, for instance, only using encryptions in which a particular selection’s hash is always numerically first within a contest. However, a malicious wrapper already knows the associations between the selection hashes, the short codes, and the actual selections, so it is not clear how a malicious wrapper could benefit from creating a bias.)

4.3 The Ballot Recording Tool

The ballot recording tool receives an election manifest, an identifier for a ballot style, the decrypted primary nonce ξ and, for a cast ballot, all the selections made by the voter. The recording tool uses the primary nonce ξ to regenerate all of the encryptions on the ballot. For a cast ballot, the tool then isolates the pre-encryptions corresponding to the selections made by the voter and, using the encryption nonces derived from the primary nonce, generates proofs of ballot-correctness as in standard ElectionGuard section 3.3.5.

Note that if a contest selection limit is greater than one, the recording tool homomorphically combines the selected pre-encryption vectors corresponding to the selections made to produce a single vector of encrypted selections. The selected pre-encryption vectors are combined by componentwise multiplication (modulo p), and the derived encryption nonces are added (modulo q) to create suitable nonces for this combined pre-encryption vector. These derived nonces will be necessary to form zero-knowledge proofs that the associated encryption vectors are well-formed.

For each uncast (implicitly or explicitly challenged) ballot, the recording tool returns the primary nonce that enables the encryptions to be opened and checked.

4.3.1 Using the Recording Tool

A wrapper for the recording tool takes one or more encrypted nonces and obtains the decryption(s) by interacting with guardians, an administrator, or a local database and then calls the recording tool with each decrypted nonce and ballot style – and for a cast ballot, the voter selections.

If the ballot is a cast ballot, the wrapper then uses the hash trimming function Ω to compute the short codes for the selections made by the voter and posts in the election record the full set of selection hashes generated from *all* pre-encryption vectors (whether or not selected by the voter), the full pre-encryption vectors corresponding to the voter selections, the proofs that these selected pre-encryption vectors are well-formed, and the short codes for the selections made by the voter.

For an uncast ballot, the wrapper computes the short codes for all possible selections and posts in the election record the full set of pre-encryption vectors, selection hashes, and short codes for each possible selection.

4.4 The Election Record

Selection vectors generated from pre-encrypted ballots are indistinguishable from those produced by standard ElectionGuard. However, the election record for each pre-encrypted ballot includes a significant amount of additional information. Specifically, for each cast ballot, the election record should contain

- the standard ElectionGuard encrypted ballot data consisting of selection vectors for each contest together with all the standard associated zero-knowledge proofs that the ballot is well-formed,
- the selection hashes for every option on the ballot (including null options) – sorted numerically within each contest, and
- the short codes and pre-encryption selection vectors associated with all selectable options (including null options) on the ballot made by the voter.

Note that in a contest with a selection limit of one, the selection vector will be identical to one of the pre-encryption selection vectors. However, when a contest has a selection limit greater than one, the resulting selection vector will be a product of multiple pre-encryption selection vectors.

For each uncast ballot, the primary nonce for that ballot is published in the encryption record.

While the basic pre-encrypted ballots are identical to standard ElectionGuard ballots, the confirmation codes are computed differently. Unlike standard ElectionGuard ballots, pre-encrypted ballot confirmation codes are computed before any selections are made. The confirmation codes on pre-encrypted ballots are computed from the full set of pre-encryptions. This is the same whether the pre-encrypted ballot is cast or spoiled.

4.4.1 Election Record Presentation

The presentation of data in the election record should be cognizant of the fact that there are two very different uses that may be made of this record. Individual voters will want to look up their own cast and uncast ballots and to easily review that they match their expectations. Election verifiers will want to verify the cryptographic artifacts associated with individual cast and uncast ballots and check their consistency as well as the consistency of the reported tallies.

It should therefore be possible for voters to see, in as clean a presentation as is feasible, the short codes associated with selections made on their cast ballots and the short codes associated with all possible selections on uncast ballots. The presentation of uncast ballots should match, as closely as feasible, the appearance of the physical uncast ballot as this presentation would facilitate comparison between the two.

For election verifiers, all of the cryptographic artifacts should be made available for verification. Verifiers should not only confirm the consistency of this additional data but also that this additional data is consistent with the cleaner data views made available to individual voters.

4.5 Verification of Pre-Encrypted Ballots

Every step of verification that applies to traditional ElectionGuard ballots also applies to pre-encrypted ballots – with the exception of the process for computing confirmation codes. However,

there are some additional verification steps that must be applied to pre-encrypted ballots. Specifically, the following verifications should be done for every pre-encrypted cast ballot contained in the election record.

- The ballot confirmation code correctly matches the hash of all contest hashes on the ballot (listed sequentially).
- Each contest hash correctly matches the hash of all selection hashes (including null selection hashes) within that contest (sorted within each contest).
- All short codes shown to voters are correctly computed from selection hashes in the election record which are, in turn, correctly computed from the pre-encryption vectors published in the election record.
- For contests with selection limit greater than 1, the selection vectors published in the election record match the product of the pre-encryptions associated with the short codes listed as selected.

The following verifications should be done for every pre-encrypted ballot listed in the election record as uncast.

- The ballot confirmation code correctly matches the hash of all contest hashes on the ballot (listed sequentially).
- Each contest hash correctly matches the hash of all selection hashes (including null selection hashes) within that contest (sorted within each contest).
- All short codes on the ballot are correctly computed from the selection hashes in the election record which are, in turn, correctly computed from the pre-encryption vectors published in the election record.
- The decryptions of all pre-encryptions correspond to the plaintext values indicated in the election manifest.

Verification 15 (Well-formedness of selection encryptions in pre-encrypted ballots)

For each possible selection on each cast ballot, an election verifier must compute the values

- (15.1) $a_0 = g^{v_0} \cdot \alpha^{c_0} \bmod p$,
- (15.2) $a_1 = g^{v_1} \cdot \alpha^{c_1} \bmod p$,
- (15.3) $b_0 = K^{v_0} \cdot \beta^{c_0} \bmod p$,
- (15.4) $b_1 = K^{w_1} \cdot \beta^{c_1} \bmod p$, where $w_1 = (v_1 - c_1) \bmod q$,
- (15.5) $c = H(\text{HE}; 0x21, K, \alpha, \beta, a_0, b_0, a_1, b_1)$.

An election verifier must then confirm the following:

- (15.A) The given values α and β are in the set \mathbb{Z}_p^* . (A value x is in \mathbb{Z}_p^* if and only if x is an integer such that $0 \leq x < p$ and $x^q \bmod p = 1$ is satisfied.)
- (15.B) The given values c_0 and c_1 are each in the set $\{0, 1, \dots, 2^{256} - 1\}$.
- (15.C) The given values v_0 and v_1 are each in the set \mathbb{Z}_q . (A value x is in \mathbb{Z}_q if and only if x is an integer such that $0 \leq x < q$.)
- (15.D) The equation $c = (c_0 + c_1) \bmod q$ is satisfied.

Verification 16 (Adherence to vote limits in pre-encrypted ballots)

For each contest on each cast ballot, an election verifier must compute the contest totals

$$(16.1) \quad \bar{\alpha} = \prod_i \alpha_i \bmod p,$$

$$(16.2) \quad \bar{\beta} = \prod_i \beta_i \bmod p,$$

where the (α_i, β_i) represent all possible selections for the contest, as well as the values

$$(16.3) \quad a_j = g^{v_j} \cdot \bar{\alpha}^{c_j} \bmod p \text{ for all } 0 \leq j \leq L,$$

$$(16.4) \quad b_j = K^{w_j} \cdot \bar{\beta}^{c_j} \bmod p, \text{ where } w_j = (v_j - jc_j) \bmod q \text{ for all } 0 \leq j \leq L,$$

$$(16.5) \quad c = H(\text{H}_E; \text{0x21}, K, \bar{\alpha}, \bar{\beta}, a_0, b_0, a_1, b_1, \dots, a_L, b_L).$$

An election verifier must then confirm the following:

(16.A) The given values c_j each satisfy $0 \leq c_j < 2^{256}$ for all $0 \leq j \leq L$.

(16.B) The given values v_j are each in \mathbb{Z}_q for all $0 \leq j \leq L$.

(16.C) The equation $c = c_0 + c_1 + \dots + c_L \bmod q$ is satisfied.

Verification 17 (Validation of confirmation codes in pre-encrypted ballots)

An election verifier must confirm the following for each pre-encrypted ballot B .

(17.A) For each selection in each contest on the ballot and the corresponding selection vector $\Psi_{i,m} = \langle E_1, E_2, \dots, E_m \rangle$ consisting of the selection encryptions $E_j = (\alpha_j, \beta_j)$, the selection hash ψ_i satisfies

$$\psi_i = H(\text{H}_E; \text{0x40}, K, \alpha_1, \beta_1, \alpha_2, \beta_2, \dots, \alpha_m, \beta_m).$$

(17.B) The contest hash χ_l for the contest with context index l for all $1 \leq l \leq m_B$ has been correctly computed from the selection hashes ψ_i as

$$\chi_l = H(\text{H}_E; \text{0x41}, l, K, \psi_{\sigma(1)}, \psi_{\sigma(2)}, \dots, \psi_{\sigma(m+L)}),$$

where σ is a permutation and $\psi_{\sigma(1)} < \psi_{\sigma(2)} < \dots < \psi_{\sigma(m+L)}$.

(17.C) The ballot confirmation code $H(B)$ has been correctly computed from the (sequentially ordered) contest hashes and if specified in the election manifest file from the additional byte array B_{aux} as

$$H(B) = H(\text{H}_E; \text{0x42}, \chi_1, \chi_2, \dots, \chi_{m_B}, B_{\text{aux}}).$$

An election verifier must also verify the following.

(17.D) There are no duplicate confirmation codes, i.e. among the set of submitted (cast and challenged) ballots, no two have the same confirmation code.

Verification 18 (Validation of short codes in pre-encrypted ballots)

An election verifier must confirm for every selectable option on every pre-encrypted ballot in the election record that the short code ω displayed with the selectable option satisfies

(18.A) $\omega = \Omega(\psi)$ where ψ is the selection hash associated with the selectable option.

Specifically, for cast ballots, this includes all short codes that are published in the election record whose associated selection hashes correspond to selection vectors that are accumulated to form tallies. For spoiled ballots, this includes all selection vectors on the ballot.

An election verifier must also confirm that for contests with selection limit greater than 1, the selection vectors published in the election record match the product of the pre-encryptions associated with the short codes listed as selected.

4.6 Hash-Trimming Functions

To allow vendors and jurisdictions to present distinct formats to their voters, the details of the hash-trimming function that produces short codes from full-sized hashes are not explicitly provided. However, to facilitate verification, a variety of possible hash-trimming functions are pre-specified here.

Two Hex Characters

- $\Omega_1(x)$ = final byte of x expressed as two hexadecimal characters

Four Hex Characters

- $\Omega_2(x)$ = final two bytes of x expressed as four hexadecimal characters

Letter-Digit

- $\Omega_3(x)$ = final byte of x expressed as a letter followed by a digit with $0, 1, \dots, 255$ mapping to A0,A1,...,A9,B0,B1,...B9,...,Z0,Z1,...,Z5

Digit-Letter

- $\Omega_4(x)$ = final byte of x expressed as a digit followed by a letter with $0, 1, \dots, 255$ mapping to 0A,0B,...,0Z,1A,1B,...1Z,...,9A,9B,...,9V

Number: 0-255

- $\Omega_5(x)$ = final byte of x expressed as a number with $0, 1, \dots, 255$ mapping to $0, 1, \dots, 255$ using the identity function

Number: 1-256

- $\Omega_6(x)$ = final byte of x expressed as a number with $0, 1, \dots, 255$ mapping to $1, 2, \dots, 256$ by adding 1

Number: 100-355

- $\Omega_7(x)$ = final byte of x expressed as a number with $0, 1, \dots, 255$ mapping to $100, 101, \dots, 355$ by adding 100

Number: 101-356

- $\Omega_8(x)$ = final byte of x expressed as a number with $0, 1, \dots, 255$ mapping to $101, 102, \dots, 355$ by adding 101

5 Hash Computation

The function H that is used throughout ElectionGuard is instantiated based on the hashed message authentication code HMAC. This section defines how to evaluate H . It first defines how inputs are represented as byte arrays and then how these inputs are used to compute hash values.

5.1 Input data representation

All inputs to the function H are byte arrays. A *byte* is a non-negative integer less than 2^8 , i.e. an integer in the set $\mathcal{B} = \{0, 1, \dots, 255\}$. Its binary form consists of at most 8 bits. Its hexadecimal form consists of at most two hexadecimal characters. Here, the leading 0 characters are written out such that a byte always has exactly two hexadecimal characters. Therefore, \mathcal{B} is represented as $\{0x00, 0x01, 0x02, \dots, 0xFE, 0xFF\}$.

A *byte array* B of length m is an array of m bytes $b_0, b_1, \dots, b_m \in \mathcal{B}$. It is represented by the concatenation³⁷ of the byte values as $B = b_0 \parallel b_1 \parallel \dots \parallel b_m$. A byte array of length m consists of $8m$ bits. For $0 \leq i < 8m$, the i -th bit of the byte array B is the $(i \bmod 8)$ -th bit of the byte $b_{\lfloor i/8 \rfloor}$.³⁸ The set of all byte arrays of length exactly m is denoted by \mathcal{B}^m and the set of all byte arrays of any length is denoted by \mathcal{B}^* .

Any byte array \mathcal{B} of length m represents a non-negative (i.e. unsigned) integer less than 2^{8m} by interpreting the bytes of \mathcal{B} as the digits of the representation in base 2^8 with the most significant bytes to the left, i.e. in big endian format. For example, the byte array $\mathcal{B} = 0x1F \parallel 0xFF$ of length 2 represents the integer $0x1FFF$ in hexadecimal form, which corresponds to the integer $2^{13} - 1 = 8191$. In general, let $b_0 \parallel b_1 \parallel \dots \parallel b_{m-1}$ be a byte array of length m , the integer

$$b_0 \cdot 2^{(m-1) \cdot 8} + b_1 \cdot 2^{(m-2) \cdot 8} + \dots + b_{m-1} \quad (97)$$

is a non-negative integer less than 2^{8m} . Vice versa, if $0 \leq a < 2^{8m}$, then the byte array of length m representing a is given as

$$\mathbf{b}(a, m) = b_0 \parallel b_1 \parallel \dots \parallel b_{m-1}, \text{ where } b_i = \lfloor a/2^{8(m-1-i)} \rfloor \bmod 2^8. \quad (98)$$

In this document and if not specified otherwise, a byte array and big endian non-negative integers are used synonymously. However, byte arrays representing integer data types have a fixed length and must be padded with $0x00$ bytes to the left. The byte array $0x00 \parallel 0x1F \parallel 0xFF$ represents the same integer as \mathcal{B} , but has length 3 instead of 2. For the integer data types used in ElectionGuard, this is laid out in detail in the following sections.

5.1.1 Integers modulo the large prime p

Most inputs to H like public keys, vote encryptions and commitments for NIZK proofs consist of big integers modulo the large prime p , i.e. integers in the set $\mathbb{Z}_p = \{0, 1, \dots, p-1\}$. Any such element is represented in big endian format by a fixed size byte array of length exactly l_p , the length

³⁷The symbol \parallel simply denotes concatenation and does not mean that this symbol is inserted as a separator into the array in any way.

³⁸Here, $\lfloor \cdot \rfloor$ denotes the floor function, which means that the result is obtained by rounding down. For a real number x , $\lfloor x \rfloor$ is the largest integer that is not larger than x .

of the byte array representing p . If p has 4096 bits like the prime given in the standard parameters in Section 3.1.1, this length is exactly 512. The conversion from an integer to a byte array works explicitly as follows. For $a \in \mathbb{Z}_p$, the byte array $\mathbf{b}(a, l_p)$ of length l_p representing a is defined as

$$\mathbf{b}(a, l_p) = b_0 \parallel b_1 \parallel \cdots \parallel b_{l_p-1}, \quad (99)$$

where

$$b_i = \lfloor a/2^{8(l_p-1-i)} \rfloor \bmod 2^8 \in \mathcal{B} \text{ for } i \in \{0, 1, \dots, l_p - 1\}. \quad (100)$$

The bytes b_i are the coefficients of a when it is written in base 2^8 , i.e.

$$a = \sum_{i=0}^{l_p-1} b_i 2^{8(l_p-1-i)} = b_0 \cdot 2^{(l_p-1) \cdot 8} + b_1 \cdot 2^{(l_p-2) \cdot 8} + \cdots + b_{l_p-1}. \quad (101)$$

Any byte array $\mathbf{b}(a, l_p)$ that represents an integer a modulo p always has length l_p . This means that, for the standard parameters, where $l_p = 512$, we have $\mathbf{b}(a, 512) = b_0 \parallel b_1 \parallel \cdots \parallel b_{511}$, where

$$b_i = \lfloor a/2^{8(511-i)} \rfloor \bmod 2^8 \in \mathcal{B} \text{ for } i \in \{0, 1, \dots, 511\}.$$

For example, $\mathbf{b}(0, 512) = 0x0000 \dots 000000000000$ is a byte array of 512 00-bytes, $\mathbf{b}(15, 512) = 0x0000 \dots 00000000000F$, $\mathbf{b}(8572345, 512) = 0x0000 \dots 00000082CDB9$, and $\mathbf{b}(p-1, 512)$ is the array shown in Section 3.1.1 for p , but ending with $\dots \text{FFFFFFFFFFFFFE}$.

5.1.2 Integers modulo the small prime q

Other inputs are integers modulo the smaller prime q , such as response values in NIZK proofs and encryption nonces. They are integers in the set $\mathbb{Z}_q = \{0, 1, \dots, q-1\}$ and are represented by a fixed size byte array of length exactly l_q in big endian format. For the standard parameters, $l_q = 32$. The conversion from integer to byte array works as above. If $a \in \mathbb{Z}_q$, the byte array $\mathbf{b}(a, 32)$ representing a is defined as

$$\mathbf{b}(a, 32) = b_0 \parallel b_1 \parallel \cdots \parallel b_{31}, \quad (102)$$

where

$$b_i = \lfloor a/2^{8(31-i)} \rfloor \bmod 2^8 \in \mathcal{B} \text{ for } i \in \{0, 1, \dots, 31\}. \quad (103)$$

Again, the bytes are coefficients of a in its 2^8 -adic form, namely

$$a = \sum_{i=0}^{31} b_i 2^{8(31-i)} = b_0 \cdot 2^{31 \cdot 8} + b_1 \cdot 2^{30 \cdot 8} + \cdots + b_{31}. \quad (104)$$

All byte arrays that represent integers modulo q have length 32. For example,

$$\begin{aligned} \mathbf{b}(0, 32) &= 0x00, \\ \mathbf{b}(16, 32) &= 0x0010, \\ \mathbf{b}(q-1, 32) &= 0xFF42. \end{aligned}$$

5.1.3 Small integers

Other integers such as *indices* are much smaller and are encoded as fixed length byte arrays in big endian format in the same way, but have much smaller lengths. In `ElectionGuard`, all such small integers are assumed to be smaller than 2^{31} . They can therefore be encoded with 4 bytes, i.e. with 32 bits and the most significant bit set to 0.³⁹

The number of guardians n and the quorum threshold value k are examples of such small integers that require even less than 4 bytes for all reasonable use cases, i.e. they are represented as $\mathbf{b}(n, 4)$ and $\mathbf{b}(k, 4)$. For example, $\mathbf{b}(5, 4) = 0x00000005$ and $\mathbf{b}(3, 4) = 0x00000003$.

5.1.4 Strings

When an input to the function H is a string s , it is encoded as a byte array $\mathbf{b}(s, \text{len}(s))$ using UTF-8 encoding. Here, $\text{len}(s)$ is the length of the UTF-8 encoding of the string s in bytes. For example, the string “ElectionGuard” is encoded as $\mathbf{b}(\text{“ElectionGuard”}, 13) = 0x456C656374696F6E4775617264$. Some string inputs may be of variable length that cannot be specified in advance by this document. Therefore all encodings of strings that are input to the `ElectionGuard` hash function H start with a 4-byte encoding in big endian format of the byte length of the string encoding followed by the encoding itself.⁴⁰ For example, if the string “ElectionGuard” is an input to H , the input byte array is $\mathbf{b}(13, 4) \parallel \mathbf{b}(\text{“ElectionGuard”}, 13) = 0x0000000D456C656374696F6E4775617264$.

5.1.5 Files

When an input to the function H is a file `file`, it is parsed as input entirely, meaning that all bytes of the file are parsed to H as a byte array $\mathbf{b}(\text{file}, \text{len}(\text{file}))$ that is a concatenation of all the bytes of the file in order. As above for strings, file lengths are not specified in this document and file inputs to H start with a 4-byte encoding⁴¹ in big endian format of the file byte length $\text{len}(\text{file})$ followed by the file bytes, i.e. $\mathbf{b}(\text{len}(\text{file}), 4) \parallel \mathbf{b}(\text{file}, \text{len}(\text{file}))$.

5.2 Hash function

The hash function H used in `ElectionGuard` is HMAC-SHA-256, i.e. HMAC⁴² instantiated with SHA-256⁴³. Therefore, H takes two byte arrays as inputs.

The first input corresponds to the key in HMAC. The HMAC-SHA-256 specification allows arbitrarily long keys, but preprocesses the key to make it exactly 64 bytes (512 bits) long, which is the block size for the input to the SHA-256 hash function. If the given key is smaller, HMAC-SHA-256 pads it with 00 bytes at the end, if it is longer, HMAC-SHA-256 hashes it first with SHA-256 and

³⁹Setting the most significant bit to 0 is done in consideration of languages and runtimes that do not have full support for unsigned integers.

⁴⁰The maximal length of strings that can be encoded for input to H is therefore 2^{32} bytes.

⁴¹Input files to H are restricted to a length of 2^{32} bytes.

⁴²NIST (2008) The Keyed-Hash Message Authentication Code (HMAC). In: FIPS 198-1. <https://csrc.nist.gov/publications/detail/fips/198/1/final>

⁴³NIST (2015) Secure Hash Standard (SHS). In: FIPS 180-4. <https://csrc.nist.gov/publications/detail/fips/180/4/final>

then pads it to 64 bytes. In ElectionGuard, all inputs that are used as the HMAC key, i.e. all inputs to the first argument of H have a fixed length of exactly 32 bytes.

The second input can have arbitrary length and is only restricted by the maximal input length for SHA-256 and HMAC. Hence we view the function H formally as follows (understanding that HMAC implementations pad the 32-byte keys to exactly 64 bytes by appending 32 0x00 bytes):

$$H : \mathcal{B}^{32} \times \mathcal{B}^* \rightarrow \mathcal{B}^{32}, H(B_0; B_1) \mapsto \text{HMAC-SHA-256}(B_0, B_1). \quad (105)$$

ElectionGuard uses HMAC not as a keyed hash function with a secret key or a message authentication code, but instead uses it as a general purpose hash function to implement a random oracle.⁴⁴

The first input is used to bind hash values to a specific election by including the parameter base hash H_P , the election base hash H_B or the extended base hash H_E . The second input consists of domain separation tags for different use cases and the actual data that is being hashed.

5.3 Hashing multiple inputs

ElectionGuard often requires that multiple input elements are hashed together. In fact, the second input to the function H specified in the previous section always consists of multiple parts. The input byte array B_1 is then simply the concatenation of the byte arrays that represent the multiple input elements as specified in Section 5.1 in the order specified in each case. As all byte arrays that represent input elements have a fixed length, there is no need for a separator byte or character.

The notation in this document lists the multiple input elements in order separated by commas. The first input element forming the byte array B_0 and the list of elements forming the second byte array B_1 are separated by a semicolon. To illustrate the notation, consider, for example, the challenge value for the proof of correctness in ballot encryption. It is computed as $H(H_E; 0x21, K, \alpha, \beta, a_0, b_0, a_1, b_1)$ (see Equation (28)). This notation means that HMAC-SHA-256 is called as HMAC-SHA-256(B_0, B_1) with the byte arrays⁴⁵

$$B_0 = H_E, B_1 = 0x21 \parallel \mathbf{b}(K, l_p) \parallel \mathbf{b}(\alpha, l_p) \parallel \mathbf{b}(\beta, l_p) \parallel \mathbf{b}(a_0, l_p) \parallel \mathbf{b}(b_0, l_p) \parallel \mathbf{b}(a_1, l_p) \parallel \mathbf{b}(b_1, l_p)$$

as inputs. The first argument is simply the extended base hash H_E . The second argument is the byte array that is the result of concatenating the domain separator byte 0x21 that identifies the use case and the byte arrays representing the public key K , the vote encryption components α and β and the commitments to the Chaum-Pedersen proofs a_0, b_0, a_1 , and b_1 in that order.

5.4 Hash function outputs

The output of SHA-256 and therefore H is a 256-bit string, which can be interpreted as a byte array of 32 bytes. As such, outputs of H are directly used as inputs to H again without any modifications.

⁴⁴Dodis Y., Ristenpart T., Steinberger J., Tessaro S. (2012) *To Hash or Not to Hash Again? (In)differentiability Results for H^2 and HMAC*. This paper shows that HMAC used as a general purpose hash function with keys of a fixed length shorter than $d - 1$, where d is the block length in bits of the underlying hash function is indistinguishable from a random oracle.

⁴⁵The symbol \parallel does not represent a separator symbol and is simply used to denote concatenation like for byte arrays as mentioned above.

When generating NIZK proofs, hash function outputs are involved in computations modulo q . Therefore they need to be interpreted as integers. The byte array $b_0 \parallel b_1 \parallel \dots \parallel b_{31}$ is interpreted as a big endian base- 2^8 representation of an integer a , which means that

$$a = \sum_{i=0}^{31} b_i 2^{8(31-i)} = b_0 \cdot 2^{31 \cdot 8} + b_1 \cdot 2^{30 \cdot 8} + \dots + b_{31}. \quad (106)$$

Therefore, hash function outputs correspond to integers in the interval $[0, 2^{256} - 1]$. Note that, although this byte representation is identical to the byte representation of integers modulo q as described above, the integers corresponding to hash function outputs do not have to be smaller than $q = 2^{256} - 189$. However, in practice, it is highly unlikely that an integer in the interval $[2^{256} - 189, 2^{256} - 1]$ will occur. When picking a 256-bit integer uniformly at random, the probability that it is not smaller than q is negligibly small, namely $189/2^{256} < 2^{-248}$.

Although hash function outputs are used in computations modulo q , they are not reduced modulo q by default. Each hash function output, such as a challenge value c in the computation of a NIZK proof, remains a byte array of length 32 that can be interpreted as an integer in $[0, 2^{256} - 1]$. Only during the computation of response values does a reduction modulo q occur.

5.5 Domain separation

The tables below list every use of the hash function H in this specification together with the specific domain separation input for the second argument. They explicitly provide the byte arrays B_0 and B_1 that are passed to HMAC as inputs to evaluate $\text{HMAC-SHA-256}(B_0, B_1)$ together with the length $\text{len}(B_1)$ of B_1 in bytes (note that always $\text{len}(B_0) = 32$). The tables use the standard parameters such that $l_p = 512$ and $l_q = 32$. For other parameters, the values l_p and l_q can be different.

5.5.1 Parameter base, manifest and base hashes

Computation of the parameter base hash H_P $H_P = H(\text{ver}; 0x00, p, q, g)$ $B_0 = \text{ver} = 0x322E302E30 \parallel b(0, 27)$ $B_1 = 0x00 \parallel b(p, 512) \parallel b(q, 32) \parallel b(g, 512)$ $\text{len}(B_1) = 1057$	§3.1.2 (4)
Computation of the manifest hash H_M $H_M = H(H_P; 0x01, \text{manifest})$ $B_0 = H_P$ $B_1 = 0x01 \parallel b(\text{len}(\text{manifest}), 4) \parallel b(\text{manifest}, \text{len}(\text{manifest}))$ $\text{len}(B_1) = 5 + \text{len}(\text{manifest})$	§3.1.3 (6)
Computation of the base hash H_B $H_B = H(H_P; 0x02, H_M, n, k)$ $B_0 = H_P$ $B_1 = 0x02 \parallel H_M \parallel b(n, 4) \parallel b(k, 4)$ $\text{len}(B_1) = 41$	§3.1.5 (7)

5.5.2 Key generation and extended base hash

NIZK proof of knowledge of polynomial coefficients in key generation $c_{i,j} = H(H_P; 0x10, i, j, K_{i,j}, h_{i,j})$ $B_0 = H_P$ $B_1 = 0x10 \parallel \mathbf{b}(i, 4) \parallel \mathbf{b}(j, 4) \parallel \mathbf{b}(K_{i,j}, 512) \parallel \mathbf{b}(h_{i,j}, 512)$ $\text{len}(B_1) = 1033$	§3.2.2 (12)
Encryption of key shares in key generation $k_{i,\ell} = H(H_P; 0x11, i, \ell, K_\ell, \alpha_{i,\ell}, \beta_{i,\ell})$ $B_0 = H_P$ $B_1 = 0x11 \parallel \mathbf{b}(i, 4) \parallel \mathbf{b}(\ell, 4) \parallel \mathbf{b}(K_\ell, 512) \parallel \mathbf{b}(\alpha_{i,\ell}, 512) \parallel \mathbf{b}(\beta_{i,\ell}, 512)$ $\text{len}(B_1) = 1545$	§3.2.2 (15)
Computation of the extended base hash H_E $H_E = H(H_B; 0x12, K)$ $B_0 = H_B$ $B_1 = 0x12 \parallel \mathbf{b}(K, 512)$ $\text{len}(B_1) = 513$	§3.2.2 (23)

5.5.3 Ballot encryption and confirmation codes

Computation of encryption nonces	§3.3.2
$\xi_{i,j} = H(H_E; 0x20, \xi_B, \text{ind}_c(\Lambda_i), \text{ind}_o(\lambda_j))$	(25)
$B_0 = H_E$	
$B_1 = 0x20 \parallel \xi_B \parallel \mathbf{b}(\text{ind}_c(\Lambda_i), 4) \parallel \mathbf{b}(\text{ind}_o(\lambda_j), 4)$	
$\text{len}(B_1) = 41$	
Computation of encryption nonces for contest data	§3.3.6
$\xi = H(H_E; 0x20, \xi_B, \text{ind}_c(\Lambda), \text{"contest_data"})$	(50)
$B_0 = H_E$	
$B_1 = 0x20 \parallel \xi_B \parallel \mathbf{b}(\text{ind}_c(\Lambda), 4) \parallel 0x0000000C636F6E7465737445F64617461$	
$\text{len}(B_1) = 53$	
Challenge computation for 0/1-range proofs	§3.3.5
$c = H(H_E; 0x21, K, \alpha, \beta, a_0, b_0, a_1, b_1)$	(28), (37)
$B_0 = H_E$	
$B_1 = 0x21 \parallel \mathbf{b}(K, 512) \parallel \mathbf{b}(\alpha, 512) \parallel \mathbf{b}(\beta, 512) \parallel \mathbf{b}(a_0, 512) \parallel \mathbf{b}(b_0, 512) \parallel \mathbf{b}(a_1, 512) \parallel \mathbf{b}(b_1, 512)$	
$\text{len}(B_1) = 3585$	
Challenge computation for L -range proofs	§3.3.8
$c = H(H_E; 0x21, K, \bar{\alpha}, \bar{\beta}, a_0, b_0, a_1, b_1, \dots, a_L, b_L)$	(56)
$B_0 = H_E$	
$B_1 = 0x21 \parallel \mathbf{b}(K, 512) \parallel \mathbf{b}(\alpha, 512) \parallel \mathbf{b}(\beta, 512) \parallel \mathbf{b}(a_0, 512) \parallel \mathbf{b}(b_0, 512) \parallel \dots \parallel \mathbf{b}(a_L, 512) \parallel \mathbf{b}(b_L, 512)$	
$\text{len}(B_1) = 1 + (2L + 5) \cdot 512$	
Key derivation for contest data encryption	§3.3.6
$k = H(H_E; 0x22, K, \alpha, \beta)$	(51)
$B_0 = H_E$	
$B_1 = 0x22 \parallel \mathbf{b}(K, 512) \parallel \mathbf{b}(\alpha, 512) \parallel \mathbf{b}(\beta, 512)$	
$\text{len}(B_1) = 1537$	
Computation of contest hashes	§3.4.1
$\chi_l = H(H_E; 0x23, \text{ind}_c(\Lambda_l), K, \alpha_1, \beta_1, \alpha_2, \beta_2, \dots, \alpha_m, \beta_m)$	(57)
$B_0 = H_E$	
$B_1 = 0x23 \parallel \mathbf{b}(\text{ind}_c(\Lambda_l), 4) \parallel \mathbf{b}(K, 512) \parallel \mathbf{b}(\alpha_1, 512) \parallel \mathbf{b}(\beta_1, 512) \parallel \mathbf{b}(\alpha_2, 512) \parallel \dots$	
$\parallel \mathbf{b}(\alpha_m, 512) \parallel \mathbf{b}(\beta_m, 512)$	
$\text{len}(B_1) = 5 + (2m + 1) \cdot 512$	
Computation of confirmation codes	§3.4.2
$H(B) = H(H_E; 0x24, \chi_1, \chi_2, \dots, \chi_{m_B}, B_{\text{aux}})$	(58)
$B_0 = H_E$	
$B_1 = 0x24 \parallel \chi_1 \parallel \chi_2 \parallel \dots \parallel \chi_{m_B} \parallel \mathbf{b}(\text{len}(B_{\text{aux}}), 4) \parallel B_{\text{aux}}$	
$\text{len}(B_1) = 5 + m_B \cdot 32 + \text{len}(B_{\text{aux}})$	
Hash chain initialization for ballot chaining	
$H_0 = H(H_E; 0x24, B_{\text{aux},0})$	(59)
$B_1 = 0x24 \parallel \mathbf{b}(\text{len}(B_{\text{aux},0}), 4) \parallel B_{\text{aux},0}, \text{len}(B_1) = 5 + \text{len}(B_{\text{aux},0})$	
Hash chain closing for ballot chaining	
$\bar{H} = H(H_E; 0x24, \bar{B}_{\text{aux}})$	(61)
$B_1 = 0x24 \parallel \bar{B}_{\text{aux}} = 0x24 \parallel H(B_\ell) \parallel \mathbf{b}(\text{len}(B_{\text{aux},0}), 4) \parallel B_{\text{aux},0} \parallel 0x00000005434C4F5345$	
$\text{len}(B_1) = 46 + \text{len}(B_{\text{aux},0})$	

5.5.4 Verifiable decryption

Challenge computation for proofs of correct decryption	§3.6.3
$c = H(H_E; 0x30, K, A, B, a, b, M)$	(71)
$B_0 = H_E$	
$B_1 = 0x30 \parallel \mathbf{b}(K, 512) \parallel \mathbf{b}(A, 512) \parallel \mathbf{b}(B, 512) \parallel \mathbf{b}(a, 512) \parallel \mathbf{b}(b, 512) \parallel \mathbf{b}(M, 512)$	
$\text{len}(B_1) = 3073$	
Challenge computation for proofs of correct decryption of contest data	§3.6.4
$c = H(H_E; 0x31, K, C_0, C_1, C_2, a, b, \beta)$	(81)
$B_0 = H_E$	
$B_1 = 0x31 \parallel \mathbf{b}(K, 512) \parallel \mathbf{b}(C_0, 512) \parallel \mathbf{b}(C_1, 512) \parallel \mathbf{b}(C_2, 512) \parallel \mathbf{b}(a, 512) \parallel \mathbf{b}(b, 512) \parallel \mathbf{b}(\beta, 512)$	
$\text{len}(B_1) = 3585$	

5.5.5 Pre-encrypted ballots

Computation of selection hashes for pre-encrypted ballots	§4.1.1
$\psi_i = H(H_E; 0x40, K, \alpha_1, \beta_1, \alpha_2, \beta_2, \dots, \alpha_m, \beta_m)$	(92)
$\psi_{m+\ell} = H(H_E; 0x40, K, \alpha_1, \beta_1, \alpha_2, \beta_2, \dots, \alpha_m, \beta_m)$	(93)
$B_0 = H_E$	
$B_1 = 0x40 \parallel \mathbf{b}(K, 512) \parallel \mathbf{b}(\alpha_1, 512) \parallel \mathbf{b}(\beta_1, 512) \parallel \mathbf{b}(\alpha_2, 512) \parallel \dots \parallel \mathbf{b}(\alpha_m, 512) \parallel \mathbf{b}(\beta_m, 512)$	
$\text{len}(B_1) = 1 + (2m + 1) \cdot 512$	
Computation of contest hashes for pre-encrypted ballots	§4.1.2
$\chi_l = H(H_E; 0x41, \text{ind}_c(\Lambda_l), K, \psi_{\sigma(1)}, \psi_{\sigma(2)}, \dots, \psi_{\sigma(m+L)})$	(94)
$B_0 = H_E$	
$B_1 = 0x41 \parallel \mathbf{b}(\text{ind}_c(\Lambda_l), 4) \parallel \mathbf{b}(K, 512) \parallel \psi_{\sigma(1)} \parallel \psi_{\sigma(2)} \parallel \dots \parallel \psi_{\sigma(m+L)}$	
$\text{len}(B_1) = 517 + (m + L) \cdot 32$	
Computation of ballot hashes for pre-encrypted ballots	§4.1.3
$H(B) = H(H_E; 0x42, K, \chi_1, \chi_2, \dots, \chi_{m_B})$	(95)
$B_0 = H_E$	
$B_1 = 0x42 \parallel \mathbf{b}(K, 512) \parallel \chi_1 \parallel \chi_2 \parallel \dots \parallel \chi_{m_B}$	
$\text{len}(B_1) = 513 + m_B \cdot 32$	
Computation of encryption nonces	§4.2.1
$\xi_{i,j,k} = H(H_E; 0x43, \xi, \text{ind}_c(\Lambda_i), \text{ind}_o(\lambda_j), \text{ind}_o(\lambda_k))$	(96)
$B_0 = H_E$	
$B_1 = 0x43 \parallel \xi \parallel \mathbf{b}(\text{ind}_c(\Lambda_i), 4) \parallel \mathbf{b}(\text{ind}_o(\lambda_j), 4) \parallel \mathbf{b}(\text{ind}_o(\lambda_k), 4)$	
$\text{len}(B_1) = 45$	

6 Verifier Construction

While it is desirable for anyone who may construct an `ElectionGuard` verifier to have as complete an understanding as possible of the `ElectionGuard` design, this section isolates the items which must be verified and maps the variables used in the specification equations herein to the labels provided in the artifacts produced in an election.

6.1 Implementation Details

There are four operations which must be performed – all on very large integer values: modular addition, modular multiplication, modular exponentiation, and hash computations, i.e. evaluations of the hash function H . These operations can be performed using a programming language that provides native support, by importing tools to perform these large integer operations, or by implementing these operations from scratch.

Modular addition

To compute $(a+b) \bmod n$, one can compute $((a \bmod n) + (b \bmod n)) \bmod n$. However, this is rarely beneficial. If it is known that $a, b \in \mathbb{Z}_n$, then one can choose to avoid the division normally inherent in the modular reduction and just use $(a+b) \bmod n = a+b$ (if $a+b < n$) or $a+b-n$ (if $a+b \geq n$).

Modular multiplication

To compute $(a \cdot b) \bmod n$, one can compute $((a \bmod n) \cdot (b \bmod n)) \bmod n$. Unless it is already known that $a, b \in \mathbb{Z}_n$, it is usually beneficial to perform modular reduction on these intermediate values before computing the product. However, it is still necessary to perform modular reduction on the result of the multiplication.

Modular exponentiation

To compute $a^b \bmod n$, one can compute $(a \bmod n)^b \bmod n$, but one should not perform a modular reduction on the exponent.⁴⁶ One should, however, never simply attempt to compute the exponentiation a^b before performing a modular reduction as the number a^b would likely contain more digits than there are particles in the universe. Instead, one should use a special-purpose modular exponentiation method such as a repeated square-and-multiply algorithm which prevents intermediate values from growing excessively large. Some languages include a native modular exponentiation primitive, but when this is not available a specialized modular exponentiation tool can be imported or written from scratch.

⁴⁶In general, if n is prime, one can compute $a^b \bmod n$ as $(a \bmod n)^{(b \bmod (n-1))} \bmod n$. In the particular instance of this specification, if $a \in \mathbb{Z}_p^*$, then one can compute $a^b \bmod p$ as $a^{(b \bmod q)} \bmod p$.

6.2 Validation Steps

This section collects the steps to validate an election and, for each step, lists the involved variables together with an explanation and pointers to how they are introduced in the specification.

6.2.1 Parameter verification

Verification 1 (Parameter validation)

An ElectionGuard election verifier must verify that it uses the correct version of the ElectionGuard specification, that it uses the standard baseline parameters, which may be hardcoded, and that the base hash values have been computed correctly.

- (1.A) The ElectionGuard specification version used to generate the election record is the same as the ElectionGuard specification version used to verify the election record.
- (1.B) The large prime is equal to the large modulus p defined in Section 3.1.1.
- (1.C) The small prime is equal to the prime q defined in Section 3.1.1.
- (1.D) The cofactor is equal to the value r defined in Section 3.1.1.
- (1.E) The generator is equal to the generator g defined in Section 3.1.1.
- (1.F) The parameter base hash has been computed correctly as

$$H_P = H(\text{ver}; 0x00, p, q, g)$$

using the version byte array $\text{ver} = 0x76322E302E30 \parallel \mathbf{b}(0, 26)$, which is the UTF-8 encoding of the version string “v2.0.0” padded with 0x00-bytes to length 32 bytes.

- (1.G) The manifest hash has been computed correctly from the manifest as

$$H_M = H(H_P; 0x01, \text{manifest}).$$

- (1.H) The base hash has been computed correctly as

$$H_B = (H_P; 0x02, H_M, n, k).$$

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
n	Total number of guardians	§3.1.2
k	Quorum of guardians required to decrypt	
H_P	Parameter base hash	
manifest	Election manifest	§3.1.3
H_M	Manifest hash	§3.1.4
H_B	Base hash	§3.1.5

6.2.2 Key ceremony verification

Verification 2 (Guardian public-key validation)

For each guardian G_i , $1 \leq i \leq n$, and for each $j \in \mathbb{Z}_k$, an election verifier must compute the value

$$(2.1) \quad h_{i,j} = g^{v_{i,j}} \cdot K_{i,j}^{c_{i,j}} \pmod p$$

and then must confirm the following.

(2.A) The value $K_{i,j}$ is in \mathbb{Z}_p^r . (A value x is in \mathbb{Z}_p^r if and only if x is an integer such that $0 \leq x < p$ and $x^q \pmod p = 1$ is satisfied.)

(2.B) The value $v_{i,j}$ is in \mathbb{Z}_q . (A value x is in \mathbb{Z}_q if and only if x is an integer such that $0 \leq x < q$.)

(2.C) The challenge $c_{i,j}$ is correctly computed as $c_{i,j} = H(\text{HP}; 0\mathbf{x}10, i, j, K_{i,j}, h_{i,j})$.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
n	Total number of guardians	§3.1.2
k	Quorum of guardians required to decrypt	
HP	Parameter base hash	
i	Sequence number of sharing guardian G_i	§3.2.2
j	Sequence number of receiving guardian G_j	
$K_{i,j}$	Public form of each random coefficient $a_{i,j}$	
$h_{i,j}$	Commitment in Schnorr proof of knowledge of coefficient $a_{i,j}$	
$c_{i,j}$	Challenge value in proof	
$v_{i,j}$	Response value in proof	

Verification 3 (Election public-key validation)

An election verifier must verify the correct computation of the joint election public key.

(3.A) The value K_i is in \mathbb{Z}_p^r and $K_i \neq 1 \pmod p$ for all $1 \leq i \leq n$.

(3.B) $K = \prod_{i=1}^n K_i \pmod p$ and $K \neq 1 \pmod p$.

variable	description	
p	4096-bit prime modulus	§3.1.1
r	Cofactor of q in $p - 1$	
n	Total number of guardians	§3.1.2
K	Joint election public key	§3.2.2
i	Sequence number of sharing guardian G_i	
K_i	Public key of guardian G_i	

6.2.3 Extended base hash verification**Verification 4 (Extended base hash validation)**

An election verifier must verify the correct computation of the extended base hash.

(4.A) $H_E = H(H_B; 0x12, K)$.

variable	description	
H_B	Base hash	§3.1.5
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3

6.2.4 Ballot verification

Verification 5 (Well-formedness of selection encryptions)

For each selectable option on each cast ballot, an election verifier must compute the values

- (5.1) $a_j = g^{v_j} \cdot \alpha^{c_j} \bmod p$ for all $0 \leq j \leq R$,
- (5.2) $b_j = K^{w_j} \cdot \beta^{c_j} \bmod p$, where $w_j = (v_j - jc_j) \bmod q$ for all $0 \leq j \leq R$,
- (5.3) $c = H(H_E; 0x21, K, \alpha, \beta, a_0, b_0, a_1, b_1, \dots, a_R, b_R)$,

where R is the option selection limit. An election verifier must then confirm the following:

- (5.A) The given values α and β are in the set \mathbb{Z}_p^r .
(A value x is in \mathbb{Z}_p^r if and only if x is an integer such that $0 \leq x < p$ and $x^q \bmod p = 1$.)
- (5.B) The given values c_j each satisfy $0 \leq c_j < 2^{256}$ for all $0 \leq j \leq R$.
- (5.C) The given values v_j are each in the set \mathbb{Z}_q for all $0 \leq j \leq R$.
(A value x is in \mathbb{Z}_q if and only if x is an integer such that $0 \leq x < q$.)
- (5.D) The equation $c = (c_0 + c_1 + \dots + c_R) \bmod q$ is satisfied.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3
(α, β)	Encryption of vote	§3.3.1
R	Option selection limit	§3.3.5
j	Index running through the range $0, 1, \dots, R$	
c_j	Derived challenge to encryption of j	
v_j	Response to challenge for encryption of j	

Verification 6 (Adherence to vote limits)

For each contest on each cast ballot, an election verifier must compute the contest totals

$$(6.1) \quad \bar{\alpha} = \prod_i \alpha_i \text{ mod } p,$$

$$(6.2) \quad \bar{\beta} = \prod_i \beta_i \text{ mod } p,$$

where the (α_i, β_i) represent all possible selections for the contest, as well as the values

$$(6.3) \quad a_j = g^{v_j} \cdot \bar{\alpha}^{c_j} \text{ mod } p \text{ for all } 0 \leq j \leq L,$$

$$(6.4) \quad b_j = K^{w_j} \cdot \bar{\beta}^{c_j} \text{ mod } p, \text{ where } w_j = (v_j - jc_j) \text{ mod } q \text{ for all } 0 \leq j \leq L,$$

$$(6.5) \quad c = H(\text{H}_E; 0\text{x}21, K, \bar{\alpha}, \bar{\beta}, a_0, b_0, a_1, b_1, \dots, a_L, b_L),$$

where L is the contest selection limit. An election verifier must then confirm the following:

(6.A) The given values α_i and β_i are each in \mathbb{Z}_p^* .

(6.B) The given values c_j each satisfy $0 \leq c_j < 2^{256}$.

(6.C) The given values v_j are each in \mathbb{Z}_q for all $0 \leq j \leq L$.

(6.D) The equation $c = (c_0 + c_1 + \dots + c_L) \text{ mod } q$ is satisfied.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3
L	Contest selection limit	§3.3.8
i	Sequence number of selections in the contest (i -th selection)	
(α_i, β_i)	Encryption of vote on i -th selection	
j	Index running through the range $0, 1, \dots, L$	
c_j	Selection limit challenge values	
v_j	Responses to selection limit challenges	

Verification 7 (Validation of confirmation codes)

An election verifier must confirm the following for each ballot B .

- (7.A) The contest hash χ_l for the contest with contest index l for all $1 \leq l \leq m_B$ has been correctly computed from the contest selection encryptions (α_i, β_i) as

$$\chi_l = H(H_E; 0x23, l, K, \alpha_1, \beta_1, \alpha_2, \beta_2 \dots, \alpha_m, \beta_m).$$

- (7.B) The ballot confirmation code $H(B)$ has been correctly computed from the contest hashes and if specified in the election manifest file from the additional byte array B_{aux} as

$$H(B) = H(H_E; 0x24, \chi_1, \chi_2, \dots, \chi_{m_B}, B_{aux}).$$

An election verifier must also verify the following.

- (7.C) There are no duplicate confirmation codes, i.e. among the set of submitted (cast and challenged) ballots, no two have the same confirmation code.

Additionally, if the election manifest file specifies a hash chain, an election verifier must confirm the following for each voting device.

- (7.D) The initial hash code H_0 satisfies $H_0 = H(H_E; 0x24, B_{aux,0})$ and $B_{aux,0}$ contains the unique voting device information.
- (7.E) For all $1 \leq j \leq \ell$, the additional input byte array used to compute $H_j = H(B_j)$ is equal to $B_{aux,j} = H(B_{j-1}) \parallel B_{aux,0}$.
- (7.F) The final additional input byte array is equal to $\bar{B}_{aux} = H(B_\ell) \parallel B_{aux,0} \parallel \mathbf{b}(\text{“CLOSE”}, 5)$ and $H(B_\ell)$ is the final confirmation code on this device.
- (7.G) The closing hash is correctly computed as $\bar{H} = H(H_E; 0x24, \bar{B}_{aux})$.

variable	description	
H_E	Extended base hash	§3.2.3
K	Joint election public key	§3.2.2
l	Sequence number of contests on the ballot	§3.4.1
Λ_l	Contest label	
i	Sequence number of selections in the contest	
(α_i, β_i)	Encryption of vote on i -th selection	
χ_l	Contest hash of the contest with contest index l	
$H(B)$	Confirmation code of ballot B	§3.4.2
B_{aux}	Additional input byte array	
$B_{\text{aux},0}$	Initial additional input byte array for ballot chaining	§3.4.3
j	Sequence number of ballots in the hash chain	
B_j	j -th ballot in the hash chain	
$B_{\text{aux},j}$	Additional input byte array for ballot chaining	
ℓ	Sequence number of last ballot in the hash chain	
\bar{B}_{aux}	Closing input byte array for ballot chaining	

6.2.5 Tally verification

Verification 8 (Correctness of ballot aggregation)

An election verifier must confirm for each option in each contest in the election manifest that the aggregate encryption (A, B) satisfies

$$(8.A) \quad A = (\prod_j \alpha_j) \bmod p,$$

$$(8.B) \quad B = (\prod_j \beta_j) \bmod p,$$

where the (α_j, β_j) are the corresponding encryptions on all cast ballots in the election record.

variable	description	
p	4096-bit prime modulus	§3.1.1
(α_j, β_j)	Encryption of selection on the j -th ballot	§3.5
(A, B)	Encrypted aggregate total of votes for this option	

Verification 9 (Correctness of tally decryptions)

For each option in each contest on each tally, an election verifier must compute the values

$$(9.1) \quad M = B \cdot T^{-1} \bmod p,$$

$$(9.2) \quad a = g^v \cdot K^c \bmod p,$$

$$(9.3) \quad b = A^v \cdot M^c \bmod p.$$

An election verifier must then confirm the following:

(9.A) The given value v is in the set \mathbb{Z}_q .

(9.B) The challenge value c satisfies $c = H(H_E; 0x30, K, A, B, a, b, M)$.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3
(A, B)	Encrypted aggregate total of votes for option	§3.5
T	Decrypted value of (A, B)	§3.6
c	Challenge to decryption	§3.6.3
v	Response to challenge	

Verification 10 (Validation of correct decryption of tallies)

An election verifier must confirm the following equations for each option in each contest in the election manifest.

(10.A) $T = K^t \text{ mod } p$.

An election verifier must also confirm that the text labels listed in the election record tallies match the corresponding text labels in the election manifest. For each contest in a decrypted tally, an election verifier must confirm the following.

(10.B) The contest text label occurs as a contest label in the list of contests in the election manifest.

(10.C) For each option in the contest, the option text label occurs as an option label for the contest in the election manifest.

(10.D) For each option text label listed for this contest in the election manifest, the option label occurs for a option in the decrypted tally contest.

An election verifier must also confirm the following.

(10.E) For each contest text label that occurs in at least one submitted ballot, that contest text label occurs in the list of contests in the corresponding tally.

variable	description	
p	4096-bit prime modulus	§3.1.1
K	Joint election public key	§3.2.2
T	Decrypted value of (A, B)	§3.6
t	Tally value	

6.2.6 Verification of correct contest data decryption

Verification 11 (Correctness of decryptions of contest data⁴⁷)

An election verifier must confirm the correct decryption of the contest data field for each contest by verifying the conditions analogous to Verification 9 for the corresponding NIZK proof with (A, B) replaced by (C_0, C_1, C_2) and M_i by m_i as follows. An election verifier must compute the following values.

$$(11.1) \quad a = g^v \cdot K^c \bmod p,$$

$$(11.2) \quad b = C_0^v \cdot \beta^c \bmod p.$$

An election verifier must then confirm the following.

(11.A) The given value v is in the set \mathbb{Z}_q .

(11.B) The challenge value c satisfies $c = H(H_E; 0x31, K, C_0, C_1, C_2, a, b, \beta)$.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3
(C_0, C_1, C_2)	Encrypted contest data	§3.6.4
β	Decryption factor of (C_0, C_1, C_2)	
c	Challenge to decryption	
v	Response to challenge	

6.2.7 Verification of challenged ballots

Verification 12 (Correctness of decryptions for challenged ballots)

For each challenged ballot, for each option in each contest on the challenged ballot, an election verifier must compute the values

$$(12.1) \quad M = \beta \cdot S^{-1} \bmod p,$$

$$(12.2) \quad a = g^v \cdot K^c \bmod p,$$

$$(12.3) \quad b = \alpha^v \cdot M^c \bmod p.$$

An election verifier must then confirm the following.

(12.A) The given value v is in the set \mathbb{Z}_q .

(12.B) The challenge value c satisfies $c = H(H_E; 0x30, K, \alpha, \beta, a, b, M)$.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3
(α, β)	Encrypted selection for option	§3.6.5
S	Decryption of (α, β)	
c	Challenge to decryption	
v	Response to challenge	

Verification 13 (Validation of correct decryption of challenged ballots)

An election verifier must confirm the correct decryption of each selection σ in each contest.

(13.A) $S = K^\sigma \text{ mod } p$.

An election verifier must also confirm that the challenged ballot is well-formed, i.e. for each contest on the challenged ballot, it must confirm the following.

(13.B) For each option in the contest, the selection σ is a valid value — usually either a 0 or a 1.

(13.C) The sum of all selections in the contest is at most the selection limit L for that contest.

An election verifier must also confirm that for each decrypted challenged ballot, the selections listed in text match the corresponding text in the election manifest.

(13.D) The contest text label occurs as a contest label in the list of contests in the election manifest.

(13.E) For each option in the contest, the option text label occurs as an option label for the contest in the election manifest.

(13.F) For each option text label listed for this contest in the election manifest, the option label occurs for a option in the decrypted challenged ballot.

variable	description	
p	4096-bit prime modulus	§3.1.1
K	Joint election public key	§3.2.2
L	Contest selection limit	§3.3.8
S	Decryption of (α, β)	§3.6.5
σ	Selection value	

Verification 14 (Correctness of contest data decryptions for challenged ballots)

An election verifier must confirm the correct decryption of the contest data field for each contest by verifying the conditions analogous to Verification 9 for the corresponding NIZK proof with (A, B) replaced by (C_0, C_1, C_2) and M by β as follows. This means it must compute the values

$$(14.1) \quad a = g^v \cdot K^c \text{ mod } p,$$

$$(14.2) \quad b = C_0^v \cdot \beta^c \text{ mod } p.$$

It must then confirm the following.

(14.A) The given value v is in the set \mathbb{Z}_q .

(14.B) The challenge value c satisfies $c = H(H_E; 0x31, K, C_0, C_1, C_2, a, b, \beta)$.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3
(C_0, C_1, C_2)	Encrypted contest data	§3.6.5
β	Decryption factor of (C_0, C_1, C_2)	
c	Challenge to decryption	
v	Response to challenge	

6.2.8 Verification of pre-encrypted ballots

Verification 15 (Well-formedness of selection encryptions in pre-encrypted ballots)

For each possible selection on each cast ballot, an election verifier must compute the values

- (15.1) $a_0 = g^{v_0} \cdot \alpha^{c_0} \bmod p$,
- (15.2) $a_1 = g^{v_1} \cdot \alpha^{c_1} \bmod p$,
- (15.3) $b_0 = K^{v_0} \cdot \beta^{c_0} \bmod p$,
- (15.4) $b_1 = K^{w_1} \cdot \beta^{c_1} \bmod p$, where $w_1 = (v_1 - c_1) \bmod q$,
- (15.5) $c = H(\mathsf{H}_E; \mathsf{0x21}, K, \alpha, \beta, a_0, b_0, a_1, b_1)$.

An election verifier must then confirm the following:

- (15.A) The given values α and β are in the set \mathbb{Z}_p^* . (A value x is in \mathbb{Z}_p^* if and only if x is an integer such that $0 \leq x < p$ and $x^q \bmod p = 1$ is satisfied.)
- (15.B) The given values c_0 and c_1 are each in the set $\{0, 1, \dots, 2^{256} - 1\}$.
- (15.C) The given values v_0 and v_1 are each in the set \mathbb{Z}_q . (A value x is in \mathbb{Z}_q if and only if x is an integer such that $0 \leq x < q$.)
- (15.D) The equation $c = (c_0 + c_1) \bmod q$ is satisfied.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3
(α, β)	Encryption of vote	§3.3.1
c_0	Derived challenge to encryption of zero	§3.3.5
c_1	Derived challenge to encryption of one	
v_0	Response to zero challenge	
v_1	Response to one challenge	

Verification 16 (Adherence to vote limits in pre-encrypted ballots)

For each contest on each cast ballot, an election verifier must compute the contest totals

$$(16.1) \quad \bar{\alpha} = \prod_i \alpha_i \bmod p,$$

$$(16.2) \quad \bar{\beta} = \prod_i \beta_i \bmod p,$$

where the (α_i, β_i) represent all possible selections for the contest, as well as the values

$$(16.3) \quad a_j = g^{v_j} \cdot \bar{\alpha}^{c_j} \bmod p \text{ for all } 0 \leq j \leq L,$$

$$(16.4) \quad b_j = K^{w_j} \cdot \bar{\beta}^{c_j} \bmod p, \text{ where } w_j = (v_j - jc_j) \bmod q \text{ for all } 0 \leq j \leq L,$$

$$(16.5) \quad c = H(H_E; 0x21, K, \bar{\alpha}, \bar{\beta}, a_0, b_0, a_1, b_1, \dots, a_L, b_L).$$

An election verifier must then confirm the following:

(16.A) The given values c_j each satisfy $0 \leq c_j < 2^{256}$ for all $0 \leq j \leq L$.

(16.B) The given values v_j are each in \mathbb{Z}_q for all $0 \leq j \leq L$.

(16.C) The equation $c = c_0 + c_1 + \dots + c_L \bmod q$ is satisfied.

variable	description	
p	4096-bit prime modulus	§3.1.1
q	256-bit prime order of subgroup of \mathbb{Z}_p^*	
r	Cofactor of q in $p - 1$	
g	Generator of subgroup of \mathbb{Z}_p^* of order q	
K	Joint election public key	§3.2.2
H_E	Extended base hash	§3.2.3
L	Contest selection limit	§3.3.8
i	Sequence number of selections in the contest (i -th selection)	
(α_i, β_i)	Encryption of vote on i -th selection	
j	Index running through the range $0, 1, \dots, L$	
c_j	Selection limit challenge values	
v_j	Responses to selection limit challenges	

Verification 17 (Validation of confirmation codes in pre-encrypted ballots)

An election verifier must confirm the following for each pre-encrypted ballot B .

- (17.A) For each selection in each contest on the ballot and the corresponding selection vector $\Psi_{i,m} = \langle E_1, E_2, \dots, E_m \rangle$ consisting of the selection encryptions $E_j = (\alpha_j, \beta_j)$, the selection hash ψ_i satisfies

$$\psi_i = H(H_E; 0x40, K, \alpha_1, \beta_1, \alpha_2, \beta_2, \dots, \alpha_m, \beta_m).$$

- (17.B) The contest hash χ_l for the contest with context index l for all $1 \leq l \leq m_B$ has been correctly computed from the selection hashes ψ_i as

$$\chi_l = H(H_E; 0x41, l, K, \psi_{\sigma(1)}, \psi_{\sigma(2)}, \dots, \psi_{\sigma(m+L)}),$$

where σ is a permutation and $\psi_{\sigma(1)} < \psi_{\sigma(2)} < \dots < \psi_{\sigma(m+L)}$.

- (17.C) The ballot confirmation code $H(B)$ has been correctly computed from the (sequentially ordered) contest hashes and if specified in the election manifest file from the additional byte array B_{aux} as

$$H(B) = H(H_E; 0x42, \chi_1, \chi_2, \dots, \chi_{m_B}, B_{aux}).$$

An election verifier must also verify the following.

- (17.D) There are no duplicate confirmation codes, i.e. among the set of submitted (cast and challenged) ballots, no two have the same confirmation code.

variable	description	
H_E	Extended base hash	§3.2.3
K	Joint election public key	§3.2.2
i	Sequence number of selections in the contest	§4.1.1
λ_i	Selection label	
(α_i, β_i)	Encryption of vote on i -th selection	
ψ_i	Selection hash	
l	Sequence number of contests on the ballot	§4.1.2
Λ_l	Contest label	
χ_l	Contest hash of the l -th contest	
$H(B)$	Confirmation code of ballot B	§4.1.3
B_{aux}	Additional input byte array	

Verification 18 (Validation of short codes in pre-encrypted ballots)

An election verifier must confirm for every selectable option on every pre-encrypted ballot in the election record that the short code ω displayed with the selectable option satisfies

(18.A) $\omega = \Omega(\psi)$ where ψ is the selection hash associated with the selectable option.

Specifically, for cast ballots, this includes all short codes that are published in the election record whose associated selection hashes correspond to selection vectors that are accumulated to form tallies. For spoiled ballots, this includes all selection vectors on the ballot.

An election verifier must also confirm that for contests with selection limit greater than 1, the selection vectors published in the election record match the product of the pre-encryptions associated with the short codes listed as selected.

variable	description	
ψ	Selection hash	§4.1.1
ω	Short code of selection hash	§4.1.4
Ω	Hash trimming function	

7 Applications to End-to-End Verifiability and Risk-Limiting Audits

The methods described in this specification can be used to enable either end-to-end (E2E) verifiability or enhanced risk-limiting audits (RLAs). In both cases, the ballots are individually encrypted and proofs are provided to allow observers to verify that the set of encrypted ballots is consistent with the announced tallies in an election.

In the case of E2E-verifiability, voters are given confirmation codes to enable them to confirm that their individual ballots are correctly recorded amongst the set of encrypted ballots. In the case of RLAs, encrypted ballots are randomly selected and compared against physical ballots to obtain confidence that the physical records match the electronic records.

To support enhanced risk-limiting audits (RLAs), it may be desirable to encrypt the ballot nonce of each ballot with a simple administrative key rather than the “heavyweight” election encryption key. This streamlines the process for decrypting an encrypted ballot that has been selected for audit. It should be noted that the privacy risks of revealing decrypted ballots are substantially reduced in the RLA case since voters are not given confirmation codes that could be used to associate them with individual ballots. The primary risk is a coercion threat (e.g., via pattern voting) that only manifests if the full set of ballots were to be decrypted.

While the administratively encrypted nonce can be stored in an electronic record alongside each encrypted ballot, one appealing RLA instantiation is for the administrative encryption of a ballot’s nonce to be printed directly onto the physical ballot. This allows an RLA to proceed by randomly selecting an encrypted ballot, fetching the associated physical ballot, extracting the nonce from its encryption on the physical ballot, using the nonce to decrypt the electronic record, and then comparing the physical ballot contents with those of the electronic record. A malicious actor with an administrative decryption key would need to go to each individual physical ballot to obtain the nonces necessary to decrypt all of the encrypted ballots, and the access to do so would enable this malicious actor to obtain all of the open ballots without necessitating the administrative decryption key.

If E2E-verifiability and enhanced RLAs are both provided in the same election, there must be separate ballot encryptions (ideally, but not necessarily, using separate election encryption keys) of each ballot. The E2E-verifiable data set must be distinguished from the enhanced RLA data set. Using the same data set for both applications would compromise voter privacy for voters whose ballots are selected for auditing.

Acknowledgments

The authors are happy to thank Eion Blanchard, Nicholas Boucher, John Caron, Joey Dodds, Felix Doerre, Gerald Doussot, Aleks Essex, Keith Fung, Rainbow Huang, Chris Jeuell, Anunay Kulshrestha, Moses Liskov, Shreyas Minocha, Arash Mirzaei, Luke Myers, Karan Newatia, Olivier Pereira, John Ramsdell, Marsh Ray, Dan Shumow, Vanessa Teague, Aaron Tomb, Daniel Wagner, Jake Waksbaum, Dan Wallach, Matt Wilhelm, and Greg Zaverucha for many helpful comments and suggestions on earlier versions of this specification.

Appendix: Other Parameters

Reduced parameters – using a 3072-bit prime

Starting from the same 256-bit prime $q = 2^{256} - 189$ for the group order, the procedure outlined in Section 3.1.1 on the standard baseline parameters can be used to find a 3072-bit prime p_{3072} . This results in the following parameters that are provided here as an alternative. The parameters satisfy all the same requirements, but the smaller prime offers faster arithmetic in the base field and thus improved performance. This comes at the cost of a somewhat lower security level. The prime p_{3072} is given as $p_{3072} = 2^{3072} - 2^{2816} + 2^{256}(\lfloor 2^{2560} \ln(2) \rfloor + \delta_{3072}) + 2^{256} - 1$ with

$$\delta_{3072} = 298707407953437995876300625370749906325322663598036756391867662926569213935809577593.$$

The hexadecimal representation of p_{3072} is as follows.

```
p3072 = 0xFFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF
      B17217F7 D1CF79AB C9E3B398 03F2F6AF 40F34326 7298B62D 8A0D175B 8BAFA2B
      E7B87620 6DEBAC98 559552FB 4AFA1B10 ED2EAE35 C1382144 27573B29 1169B825
      3E96CA16 224AE8C5 1ACBDA11 317C387E B9EA9BC3 B136603B 256FA0EC 7657F74B
      72CE87B1 9D6548CA F5DFA6BD 38303248 655FA187 2F20E3A2 DA2D97C5 0F3FD5C6
      07F4CA11 FB5BFB90 610D30F8 8FE551A2 EE569D6D FC1EFA15 7D2E23DE 1400B396
      17460775 DB8990E5 C943E732 B479CD33 CCCC4E65 9393514C 4C1A1E0B D1D6095D
      25669B33 3564A337 6A9C7F8A 5E148E82 074DB601 5CFE7AA3 0C480A54 17350D2C
      955D5179 B1E17B9D AE313CDB 6C606CB1 078F735D 1B2DB31B 5F50B518 5064C18B
      4D162DB3 B365853D 7598A195 1AE273EE 5570B6C6 8F969834 96D4E6D3 30D6E582
      CAB40D66 550984EF 0C42A457 4280B378 45189610 AE3E4BB2 2590A08F 6AD27BFB
      FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF
```

The hexadecimal representation of the cofactor $r_{3072} = (p_{3072} - 1)/q$ is shown below.

```
r3072 = 0x01 00000000 00000000 00000000 00000000 00000000 00000000 00000000
      B17217F7 D1CF79AB C9E3B398 03F2F6AF 40F34326 7298B62D 8A0D175B 8BAB857A
      E8F42816 5418806C 62B0EA36 355A3A73 E0C74198 5BF6A0E3 130179BF 2F0B43E3
      3AD86292 3861B8C9 F768C416 9519600B AD06093F 964B27E0 2D868312 31A9160D
      E48F4DA5 3D8AB5E6 9E386B69 4BEC1AE7 22D47579 249D5424 767C5C33 B9151E07
      C5C11D10 6AC446D3 30B47DB5 9D352E47 A53157DE 04461900 F6FE360D B897DF53
      16D87C94 AE71DAD0 BE84B647 C4BCF818 C23A2D4E BB53C702 A5C8062D 19F5E9B5
      033A94F7 FF732F54 12971286 9D97B8C9 6C412921 A9D86797 70F499A0 41C297CF
      F79D4C91 49EB6CAF 67B9EA3D C563D965 F3AAD137 7FF22DE9 C3E62068 DD0ED615
      1C37B4F7 4634C2BD 09DA912F D599F433 3A8D2CC0 05627DCA 37BAD43E 64CAF318
      9FD4A7F5 29FD4A7F 529FD4A7 F529FD4A 7F529FD4 A7F529FD 4A7F529F D4A7F52A
```

And the generator $g_{3072} = 2^{r_{3072}} \bmod p_{3072}$ has the following hexadecimal representation.

```
g3072 = 0x4A1523CB 0111B381 04EBCDE5 163F581E EEDD9163 7AC57544 C1D22832 34272732
```

FF0CD85F 38539544 3F573701 32A237FF 38702AB0 37F35E7C 7003669D 83697BA1
3BED69B6 3C88BD61 0D33C6A8 9E4882EE 6F849F05 06A4A8F0 B169E5CA 000A21DC
16D7DCEC C69E593C 65967739 3B6CE260 D3D6A578 E74E42A1 B2ADE1ED 8627050C
FB59E604 CAC389E9 9161DA6E 6E9407DF 94517864 01003A8B 7626AC5E 90B888EA
BB5E07E9 96B18662 9B17165F D630E139 788F674D FF4978A6 B74C6D02 0A6570CC
7C7A9E38 21283571 BA3FA1FC C6901A8C 28D02EF8 B8C4B019 F7DDADE5 1A089C57
EF90C2CE 50761754 D778BC9A BFD84809 5C4A0ED0 FA7B7AE5 2CDA4BD6 E2CB16F3
8EDC033F 32F259C5 13DD9E0D 1F780886 D71D7DB8 35F3F08D B11CC9CD 41EB0D5A
37AC6DBA 1A1EBA55 C378BC06 95B9D93A A59903EB A1CE5288 6A0BAAF8 15354863
1BCEAC52 07B97205 BE8FDF83 0F27348C 7AE852F9 F8876887 D23B8054 A077DC8A
EC0BF615 A1FA74BC 727014CF AC40E20E A194489F 63A6C224 27CB999C 9D04AA61

Toy parameters for testing purposes only

The following parameter sets are provided for testing purposes only and must not be used in any real-world scenario. The involved primes are too small to provide any security at all.

- 7-bit q , 16-bit p

$$\begin{aligned}q &= 127 = 0x007F, \\p &= 59183 = 0xE72F, \\r &= 466 = 0x01D2, \\g &= 32616 = 0x7F68\end{aligned}$$

- 16-bit q , 32-bit p

$$\begin{aligned}q &= 65521 = 0xFFFF1, \\p &= 4214179679 = 0xFB2B 475F, \\r &= 64318 = 0xFB3E, \\g &= 496451214 = 0x1D97 3E8E\end{aligned}$$

- 16-bit q , 48-bit p

$$\begin{aligned}q &= 65521 = 0xFFFF1, \\p &= 281010572049407 = 0xFF93 DF53 3BFF, \\r &= 4288862686 = 0xFFA2 D9DE, \\g &= 109132885510074 = 0x6341 7ADF C7BA\end{aligned}$$

- 24-bit q , 64-bit p

$$\begin{aligned}q &= 16777213 = 0xFFFFFFFF, \\p &= 18444843247520538623 = 0xFFF93F35 6A395FFF, \\r &= 1099398526294 = 0x000000FF F9423556, \\g &= 14757355607624201864 = 0xCCCCA8BC C08F3688\end{aligned}$$

- 32-bit q , 96-bit p

$q = 4294967291 = 0xFFFFFFFFB,$
 $p = 79227651399410325621583970303$
 $= 0xFFFF93C4\ 6882B6AA\ F57CFFFF,$
 $r = 18446625091983808922$
 $= 0xFFFF93C9\ 6880999A,$
 $g = 60786014689883675535506158858$
 $= 0xC469034B\ 2CE5EC6E\ 1970350A$

- 32-bit q , 128-bit p

$q = 4294967291 = 0xFFFFFFFFB,$
 $p = 340282366887442052436576802921059975167$
 $= 0xFFFFFFFF\ 93C46B0F\ B6C381D8\ FFFFFFFF,$
 $r = 79228162598699067120922761626$
 $= 0x00000001\ 00000004\ 93C46B26\ 9999999A,$
 $g = 55628101181055236817878380639043675517$
 $= 0x29D99524\ 0DFB12B3\ 6FD0F8CCE06B657D$

- 48-bit q , 192-bit p

$q = 0x0000FFFF\ FFFFFFFC5,$
 $p = 0xFFFFFFFF\ FFFFFFFF\ 9ECB7796\ 49D9A82D\ FFFFFFFF\ FFFFFFFF,$
 $r = 0x00010000\ 0000003A\ FFFF9ECB\ 852F49C3\ 4115B1E6,$
 $g = 0x0B5DA090\ 0B367E3C\ 92A11019\ 54DB5E3C\ 873E929A\ 0E324F00$

- 64-bit q , 256-bit p

$q = 0xFFFFFFFF\ FFFFFFFC5,$
 $p = 0xFFFFFFFF\ FFFFFFFF\ 93C467E3\ 7DB1212B\ 89995855\ 493FF059\ FFFFFFFF\ FFFFFFFF,$
 $r = 0x00000001\ 00000000\ 0000003A\ 93C467E3\ 7DB12EAB\ 97DD49C3\ 4115B1E6,$
 $g = 0x3B543166\ 9E3E4893\ DF745C67\ CDCFD95C\ CDDA2248\ 78A3CD5D\ 3226F75C\ C5A95638$

Glossary of typical variable usage

variable	description
(a, b)	Commitment to an encryption
(A, B)	Aggregation of multiple encrypted votes
$a_{i,j}$	Polynomial coefficient of secret polynomial of guardian G_i
B	Ballot
c	Challenge value
D	Contest data field
g	Generator of order- q group of \mathbb{Z}_p^*
H	Hash function
i, j	Index values
k	Decryption threshold
K	Joint election public key
K_i	Public key of i^{th} guardian
$K_{i,j}$	Public commitment of polynomial coefficient $a_{i,j}$
l, ℓ	Index values
L	Selection or range limit
M	Exponentiated decryption
M_i	Partial decryption of i^{th} guardian
n	Number of guardians
p	Large prime modulus
q	Small prime order of subgroup of \mathbb{Z}_p^*
r	Cofactor of q in $p - 1$
R	Range limit or option selection limit
s	Secret value
S	Encoded selection value
t	Tally value
T	Encoded tally value
u, v	Proof response values
w	Lagrange coefficient

variable	description
(α, β)	Encryption of individual vote
ξ	Random nonce
χ	Contest hash
σ	Selection value
ψ	Selection hash
Ψ	Selection vector
ω	Short code of selection hash
Ω	Hash trimming function