

Post-Quantum Anonymous Veto Networks

Jintai Ding¹[0000–0003–1257–7598], Doug Emery¹, Johannes Müller²[0000–0003–2134–3099], Peter Y.A. Ryan², and Vonn Kee Wong¹

¹ University of Cincinnati, Ohio, USA
jintai.ding at gmail.com

{ wongvg, emerydg } at mail.uc.edu
² SnT, University of Luxembourg, Luxembourg
{ johannes.mueller, peter.ryan } at uni.lu

Abstract. Anonymous veto networks (AV-nets), originally proposed by Hao and Zielinski (2006), are particularly lightweight protocols for evaluating a veto function in a peer-to-peer network such that anonymity of all protocol participants is preserved. Prior to this work, anonymity in all AV-nets from the literature relied on the decisional Diffie-Hellman (DDH) assumption and can thus be broken by (scalable) quantum computers. In order to defend against this threat, we propose two practical and completely lattice-based AV-nets. The first one is secure against passive and the second one is secure against active adversaries. We prove that anonymity of our AV-nets reduces to the ring learning with errors (RLWE) assumption. As such, our AV-nets are the first ones with post-quantum anonymity. We also provide performance benchmarks to demonstrate their practicality.

1 Introduction

In many jury or executive committee votings, certain results are only effective if supported by all members. Such votings, of which there are many instances in the real world, are called *veto votings*. Very recently, for example, the Supreme Court of the United States ruled that guilty verdicts for criminal trials be unanimous.³ In order to protect each voter’s freewill, veto votings are often required to not reveal any sensitive information except for the final result, i.e., whether or not at least one voter vetoed. Such votings are called *anonymous veto votings*.

Solutions for *electronic* anonymous veto protocols have a long history. In fact, David Chaum proposed the first such protocol, named *dining cryptographers network (DC-net)*, more than three decades ago [5, 6]. Since Chaum’s original protocol returns the correct result if and only if an odd number of voters decides to veto, modifications of Chaum’s protocol have been proposed to solve these and further issues (see, e.g., [11]).

However, DC-nets assume pairwise shared keys among the voters and their complexity is quadratic in the number of voters. In order to overcome these limitations, Hao and Zielinski introduced the concept of *anonymous veto networks*

³ Ramos v. Louisiana, No. 18-5925, 590 U.S. __ (2020).

(*AV-nets*) (originally proposed in [13], with some extensions in [1]). In contrast to DC-nets, AV-nets are very lightweight, both regarding the number of rounds, computation, bandwidth and system complexity.

Anonymity of existing AV-nets from the literature relies on the hardness of the decisional Diffie-Hellman (DDH) problem. Since this problem could efficiently be solved by (scalable) future quantum computers, no AV-net with *post-quantum anonymity* has been proposed prior to our work. Unfortunately, as we will explain in Section 2.3, the fact that previous AV-nets are tailored specifically to the DDH problem makes it infeasible to transform them into AV-nets with post-quantum anonymity in a straightforward way.

Our contributions. We present the first completely lattice-based AV-nets. Our protocols are efficient and practically realizable. Anonymity of voters relies on the decisional ring learning with errors (RLWE) assumption. Using the RLWE assumption in our protocol is inspired from [17, 18] in which an RLWE analogue of the Diffie-Hellman key exchange was proposed. Our protocols do not require a central tallying authority; instead the voters themselves securely compute the final result. More precisely, we provide the following contributions:

1. We propose a 2-round lattice-based AV-net that is secure against *passive* (honest-but-curious) adversaries (Section 4). We first precisely describe this protocol (Section 4.1), then show that it produces the correct final result (Section 4.2), and that anonymity/privacy of the voters is guaranteed under RLWE if all but two voters are corrupted by a passive adversary (Section 4.3).
2. We propose a 4-round lattice-based AV-net that is secure against *active* (malicious) adversaries (Section 5). We first precisely describe this protocol (Section 5.1), then show that the correctness of the final result can publicly be verified (Section 5.2), and that anonymity/privacy of the voters is guaranteed under RLWE if all but two voters are corrupted by an active adversary (Section 5.3).
3. We provide experimental performance benchmarks of our lattice-based AV-nets (Section 6).
4. We discuss the properties of the two lattice-based AV-nets as well as possible alternative approaches (Section 7).

We note that, in the remainder of this paper, we use the expressions “privacy” and “anonymity” interchangeably.

2 AV-Net by Hao and Zielinski

In this section, we first describe the original AV-net proposed by Hao and Zielinski [13] which provides anonymity under the DDH assumption. We then elaborate on why building AV-nets with lattice-based anonymity is challenging and requires careful attention.

2.1 Protocol description

The main idea behind the AV-net protocol by Hao and Zielinski [13] is the following one. The protocol is divided into an offline and an online phase. In the offline phase, the voters collaboratively generate certain related blinding elements, one individual element y_i for each voter V_i . In the subsequent online phase, voters can then decide to either veto or not. If V_i decides not to veto, then she raises y_i (as generated in the offline phase) to a specific integer s_i , and to a random integer r_i , otherwise. After that, all blinded choices are homomorphically aggregated. Furthermore, both in the offline and the online phase, zero-knowledge proofs (ZKPs) of knowledge are integrated to guarantee that voters choose their (otherwise malleable) messages pairwise independently.

The specific structure of the blinding elements y_1, \dots, y_m generated in the offline phase ensures that the result of the homomorphic aggregation equals 1 if and only if all voters choose “no veto”. The technical mechanism behind this concept is based on the following result (details will become clear further below).

Lemma 1. *Let R be a commutative ring. Let r_1, \dots, r_m be elements in R . Then the following equation holds true:*

$$\sum_{i=1}^m \sum_{j=1}^{i-1} r_i \cdot r_j = \sum_{i=1}^m \sum_{j=i}^m r_i \cdot r_j$$

Proof. See [13].

Let us now describe the AV-net protocol by Hao and Zielinski [13] with full technical details.

Protocol participants. The AV-net protocol is run among the following participants:

- Voters V_1, \dots, V_m .
- Bulletin board B .

We assume that for each voter V_i , there exists a mutually authenticated channel between V_i and the bulletin board B .

Parameters. Let G be finite cyclic group of prime order q with generator g . We assume that the decisional Diffie-Hellman (DDH) assumption holds true in G , i.e., the following two distributions are computationally indistinguishable:

- (g^a, g^b, g^{ab}) , where $a, b \xleftarrow{r} \mathbb{Z}_q$.
- (g^a, g^b, g^c) , where $a, b, c \xleftarrow{r} \mathbb{Z}_q$.

Offline phase. Each voter V_i runs the following program:

1. $s_i \xleftarrow{r} \mathbb{Z}_q$
2. $h_i \leftarrow g^{s_i}$
3. $\pi_i^1 \leftarrow \text{ZKP of knowledge of } \log_g h_i$
4. Publish (π_i^1, h_i)

After all voters have published their h_i 's (equipped with valid ZKPs), each voter V_i (locally) computes her individual blinding element y_i as follows:

$$y_i \leftarrow \left(\prod_{j=1}^{i-1} h_j \right) \cdot \left(\prod_{j=i+1}^m h_j \right)^{-1}.$$

Online phase. Voter V_i computes her “encrypted” choice as follows:

1. If “no veto”, then set $c_i \leftarrow y_i^{s_i}$.
2. If “veto”, then choose $r_i \xleftarrow{r} \mathbb{Z}_q$, and set $c_i \leftarrow y_i^{r_i}$.
3. $\pi_i^2 \leftarrow \text{ZKP of knowledge of } \log_{y_i} c_i$
4. Publish (π_i^2, c_i)

After all voters have published their c_i 's (equipped with valid ZKPs), each voter (locally) computes the final result as follows:

$$\text{res} \leftarrow \begin{cases} \text{no veto} & \text{if } \prod_{i=1}^m c_i = 1 \\ \text{veto} & \text{otherwise} \end{cases}.$$

2.2 Correctness and anonymity

We now describe why the AV-net by Hao and Zielinski is correct and provides anonymity under the DDH assumption. We focus on the case of passive adversaries; the ZKPs invoked ensure that the AV-net is also secure against active adversaries (see [13] for details).

Correctness. Let us first assume that all voters choose “no veto”. Then, we have that

$$\begin{aligned} \prod_{i=1}^m c_i &= \prod_{i=1}^m y_i^{s_i} = \prod_{i=1}^m \left(\left(\prod_{j<i} h_j \right) \left(\prod_{j>i} h_j \right)^{-1} \right)^{s_i} \\ &= g^{(\sum_{i=1}^m \sum_{j<i} s_i s_j) - (\sum_{i=1}^m \sum_{j>i} s_i s_j)} = g^0 = 1 \end{aligned}$$

holds true, where the second but last equality follows from Lemma 1. Conversely, assume that (at least) one voter vetoes, say voter V_l . Then, we have that

$$\prod_{i=1}^m c_i = y_l^{r_l} \cdot \prod_{i \neq l} c_i$$

is distributed uniformly at random in G . Hence, if $|G|$ is sufficiently large, then the probability that this product equals 1 is negligible.

Anonymity. Let V_i be an arbitrary (honest) voter. Assume that at least one further voter V_j is honest, too. Then, the sum $\sum_{j<i} s_j - \sum_{j>i} s_j$ is distributed uniformly at random in \mathbb{Z}_q . Hence, if V_i does not veto, then the triple

$$(h_i, y_i, c_i) = (g^{s_i}, g^{(\sum_{j<i} s_j) - (\sum_{j>i} s_j)}, g^{s_i \cdot ((\sum_{j<i} s_j) - (\sum_{j>i} s_j))})$$

is a DDH-triple, and otherwise a random triple

$$(h_i, y_i, c_i) = (g^{s_i}, g^{(\sum_{j<i} s_j) - (\sum_{j>i} s_j)}, g^{s_i \cdot r_i}).$$

Under the assumption that the DDH problem is intractable in G , it is not possible to distinguish between these two distributions.

2.3 Challenges for lattice-based anonymity

As we have seen in Section 2.2, the design of [13] is tailored specifically to reduce anonymity to the DDH-assumption. Therefore, if we want to design an AV-net whose anonymity reduces to a different (e.g., lattice-based) hardness assumption, then we have to adapt all technical details accordingly. This is even more challenging in the case of lattice-based anonymity: controlling the noise of lattice-based cryptographic primitives is non-trivial and requires careful attention.

Furthermore, the original AV-net [13] includes ZKPs of knowledge to defend against active adversaries which choose their messages in relation to the honest voters' ones. Even though there exist efficient lattice-based ZKPs in the literature, these ZKPs are tailored to specific lattice-based primitives. Unfortunately, it is not immediately clear how to employ these primitives to construct a lattice-based AV-net. Therefore, we decided to construct an actively secure lattice-based AV-net without ZKPs altogether (Section 5).

3 Cryptographic Primitives

In this section, we introduce the cryptographic primitives that we later employ in our lattice-based veto protocols (Section 4 and 5). Throughout this paper, we use the following parameters and conventions:

- Let n be a power of 2.
- Let R be the cyclotomic ring $\mathbb{Z}[X]/f(X)$ where $f(X) = X^n + 1$.
- Let q be a prime such that $q \equiv 1 \pmod{2n}$.
- Let R_q be the quotient ring R/qR .
- Let the coefficients of a polynomial in R_q be in the interval $[-\frac{q-1}{2}, \frac{q-1}{2}]$.
- Let $\|\cdot\|$ be the ℓ_2 -norm on R_q and $\|\cdot\|_\infty$ be the ℓ_∞ -norm on R_q .
- Let m be an integer. (This will be the number of voters.)
- Let $\Lambda = \mathbb{Z}^n$.
- Let $\rho_\sigma(\mathbf{x}) = e^{-\pi\|\mathbf{x}\|^2/\sigma^2}$ be the Gaussian function on \mathbb{R}^n with center at the zero vector and the parameter σ .

- Let $\rho_\sigma(\Lambda) = \sum_{\mathbf{x} \in \Lambda} \rho_\sigma(\mathbf{x})$ be the discrete integral of ρ_σ over Λ .
- Let $D_{\Lambda, \sigma}$ be the discrete Gaussian distribution over Λ with center at zero vector and parameter σ . For all $\mathbf{y} \in \Lambda$, we have $D_{\Lambda, \sigma}(\mathbf{y}) = \frac{\rho_\sigma(\mathbf{y})}{\rho_\sigma(\Lambda)}$.
- Let χ be the discrete Gaussian distribution $D_{\mathbb{Z}_q^n, \sigma}$.

The *decisional ring learning with errors (RLWE) problem* is about determining whether a list of polynomial pairs $(a_i, b_i) \in R_q \times R_q$ were generated uniformly at random or were constructed such that a_i is chosen uniformly at random while $b_i = a_i \cdot s_i + e_i$, where $s_i \in R_q$ is the secret and $e_i \leftarrow \chi$ is the error.

The parameters are chosen to satisfy the theorem below:

Theorem 1 ([20]). *For n, R, q and β as defined above, there is an efficiently samplable distribution χ over R with $\Pr[\|x\| > \beta : x \leftarrow \chi] \leq \text{negl}(n)$, such that if there exists an efficient algorithm that solves $RLWE_{n, q, \chi}^{(m)}$, then there is an efficient quantum algorithm for solving $n^{2.5}(q/\beta)(nm/\log(nm))^{1/4}$ -approximate worst-case SVP for ideal lattices over R .*

We recall some useful lemmas.

Lemma 2 ([16], Lemma 2.5). *For $\sigma > 0, r \geq 1/\sqrt{2\pi}$, $\Pr[\|x\| > r\sigma\sqrt{n} : x \leftarrow D_{\mathbb{Z}^n, \sigma}] < (\sqrt{2\pi}er^2 \cdot e^{-\pi r^2})^n$.*

Lemma 3 ([19], Lemma 2). *For $a, b \in R_q$, $\|a \cdot b\|_\infty \leq \|a\| \cdot \|b\|$.*

In addition, we let $\beta = r\sigma\sqrt{n}$ and we need to carefully choose $r \geq 1/\sqrt{2\pi}$ so that choosing x from $D_{\mathbb{Z}^n, \sigma}$ with ℓ_2 -norm greater than β has negligible probability according to Lemma 2.

Furthermore, to ensure the correctness of our veto protocol, we require that

$$\frac{q}{4} - 2 \geq m(m-1)\beta^2 + m\beta$$

holds true.

4 Passively Secure Lattice-Based AV-Net

The following AV-net protocol provides privacy in the presence of passive (honest-but-curious) adversaries. In Section 5, we show how to extend this AV-net such that privacy can be guaranteed even if all but two voters actively deviate from their specified programs.

In what follows, we first describe the passively secure AV-net protocol with full technical details (Section 4.1), then we prove that this protocol is correct (Section 4.2), and eventually elaborate on the privacy it provides (Section 4.3).

4.1 Protocol description

We use the same protocol participants as in the original AV-net (Section 2.1).

Parameters. We briefly recall the main parameters from Section 3 that we use in the passively secure veto protocol. Essentially, all computation is done in the ring R_q . The distribution χ samples elements from R_q such that $\text{RLWE}_{n,q,\chi}^{(m)}$ holds true. Let a be an element from R_q chosen uniformly at random. In what follows, we implicitly assume that all protocol participants take these parameters as input.

Offline phase. If, in the online phase (see below), voter V_i chooses “no veto”, then she uses a specific element $y_i \in R_q$ to blind her choice, and a random element otherwise. The elements y_1, \dots, y_m (for voters V_1, \dots, V_m) will have a specific structure such that

- the distribution of blinded “veto” choices is indistinguishable from the uniform distribution over R_q (under RLWE) which itself is the distribution of “no veto” choices, and
- all blinding elements collectively equal out if and only if all voters choose “veto”.

In fact, each voter’s blinding element y_i is a specific linear combination of elements b_j that are generated by all the other voters V_j ($i \neq j$). More precisely, each voter V_i generates b_i as follows:

1. Choose $s_i, e_i \leftarrow \chi^2$.
2. Set $b_i \leftarrow a \cdot s_i + e_i$.

After all voters have published their b_i ’s, each voter V_i (locally) computes her individual blinding element y_i as follows:

$$y_i \leftarrow \left(\sum_{j=1}^{i-1} b_j \right) - \left(\sum_{j=i+1}^m b_j \right).$$

Online phase. Voter V_i computes her “encrypted” choice as follows:

1. If “no veto”, then choose $e'_i \leftarrow \chi$, and set $c_i \leftarrow s_i y_i + e'_i$.
2. If “veto”, then choose $r_i \xleftarrow{r} R_q$, and set $c_i \leftarrow r_i$.

After all voters have published their c_i ’s, each voter (locally) computes the final result as follows:

$$\text{res} \leftarrow \begin{cases} \text{no veto} & \text{if } \left\| \sum_{i=1}^m c_i \right\|_{\infty} \leq \frac{q}{4} - 2 \\ \text{veto} & \text{otherwise} \end{cases}.$$

4.2 Correctness

In this section, we show that the veto protocol, as defined in Section 4.1, is correct, i.e., it outputs the correct result (with overwhelming probability) if all

participants follow the protocol specification correctly (Theorem 2). To this end, we use the following result which ensures that the error terms introduced (for privacy reasons) do not undermine correctness of the veto protocol except for with negligible probability.

Lemma 4. *The probability that a uniformly chosen random element $r \in R_q$ has max norm less than or equal to $N \geq 1$ is given by*

$$\Pr[\|x\|_\infty \leq N : x \leftarrow R_q] = \frac{(2N + 1)^n}{q^n}.$$

Theorem 2 (Correctness). *Let P be the veto protocol defined in Section 4.1. Assume that all voters V_1, \dots, V_m (and the bulletin board B) are honest, i.e., run their programs as specified by the protocol. Then, we have that for all runs (of this instance) of P , the following equivalence holds true with overwhelming probability: The final result res is “veto” if and only if there exists (at least) one voter V_i who chooses “veto”.*

Proof. Let us start with a variant of the veto protocol without error terms, i.e., $e_i, e'_i = 0$ for all voters V_i .⁴ Now, if all voters choose “no veto”, we have that

$$\begin{aligned} \sum_{i=1}^m c_i &= \sum_{i=1}^m s_i \cdot y_i = \sum_{i=1}^m s_i \cdot \left(\left(\sum_{j=1}^{i-1} b_j \right) - \left(\sum_{j=i+1}^m b_j \right) \right) \\ &= \sum_{i=1}^m s_i \cdot \left(a \cdot \left(\left(\sum_{j=1}^{i-1} s_j \right) - \left(\sum_{j=i+1}^m s_j \right) \right) \right) \\ &= a \cdot \left(\left(\sum_{i=1}^m \sum_{j=1}^{i-1} s_i \cdot s_j \right) - \left(\sum_{i=1}^m \sum_{j=i+1}^m s_i \cdot s_j \right) \right) = 0 \end{aligned}$$

holds true, where the last equation follows from Lemma 1.

Conversely, if (at least) one voter vetoed, then the sum $\sum_{i=1}^m c_i$ is distributed uniformly at random over R_q . Hence, res correctly reflects how voters voted in the veto protocol (without error terms).

⁴ We note that, in this case, the protocol would not guarantee privacy.

Now, let us return to the actual veto protocol, including error terms. Assuming that all voters choose “no veto”, we have that

$$\begin{aligned}
\sum_{i=1}^m c_i &= \sum_{i=1}^m (s_i \cdot y_i + e'_i) \\
&= \sum_{i=1}^m s_i \cdot \left(\left(\sum_{j=1}^{i-1} b_j \right) - \left(\sum_{j=i+1}^m b_j \right) \right) + \sum_{i=1}^m e'_i \\
&= \sum_{i=1}^m s_i \cdot \left(\left(\sum_{j=1}^{i-1} a \cdot s_j + e_j \right) - \left(\sum_{j=i+1}^m a \cdot s_j + e_j \right) \right) + \sum_{i=1}^m e'_i \\
&= \left(\sum_{i=1}^m \sum_{j=1}^{i-1} s_i \cdot e_j - \sum_{i=1}^m \sum_{j=i+1}^m s_i \cdot e_j \right) + \sum_{i=1}^m e'_i
\end{aligned}$$

holds true, where the last equality follows from what we have shown above for the variant without error terms.

Recall that all s_i, e_i and e'_i are chosen according to χ , hence their norm is bounded by β (with overwhelming probability in the security parameter n). Therefore by Lemma 3, we have that (with overwhelming probability after carefully choosing r in Lemma 2)

$$\begin{aligned}
\left\| \sum_{i=1}^m c_i \right\|_{\infty} &= \left\| \left(\sum_{i=1}^m \sum_{j=1}^{i-1} s_i \cdot e_j - \sum_{i=1}^m \sum_{j=i+1}^m s_i \cdot e_j \right) + \sum_{i=1}^m e'_i \right\|_{\infty} \\
&\leq \sum_{i=1}^m \sum_{j \neq i} \|s_i \cdot e_j\|_{\infty} + \sum_{i=1}^m \|e'_i\|_{\infty} \\
&\leq \sum_{i=1}^m \sum_{j \neq i} \|s_i\| \cdot \|e_j\| + \sum_{i=1}^m \|e'_i\| \\
&\leq m \cdot (m-1) \cdot \beta^2 + m \cdot \beta \\
&\leq \frac{q}{4} - 2.
\end{aligned}$$

Conversely, assume that one of the voters vetoes, hence chooses c_i uniformly at random from R_q . From Lemma 4, it follows that the probability that $\sum_{i=1}^m c_i$ has max norm $\leq \frac{q}{4} - 2$ is negligible:

$$\Pr[\|r\|_{\infty} \leq \frac{q}{4} - 2 : r \leftarrow R_q] = 2^{-n} \left(\frac{q-6}{q} \right)^n < 2^{-n}. \quad (1)$$

Hence, altogether, we can conclude that (with overwhelming probability) the final result res equals “veto” if and only if at least one voter vetoes. This proves the correctness of the veto protocol defined in Section 4.1.

4.3 Privacy

In this section, we show that the veto protocol, as defined in Section 4.1, provides privacy in the presence of honest-but-curious adversaries. The privacy notion we apply follows [3].

Theorem 3 (Privacy). *Assume that $RLWE_{n,q,\chi}^{(m)}$ holds true. Let A be an arbitrary passive ppt adversary which controls (at most) all but two voters $(V_i)_{i \in \mathcal{I}_{dis}}$. Let $(V_i)_{i \in \mathcal{I}_{hon}}$ denote the remaining (uncorrupted) voters. Let $(v_i)_{i \in \mathcal{I}_{hon}}$ and $(v'_i)_{i \in \mathcal{I}_{hon}}$ be two arbitrary vectors of choices that yield the same result \mathbf{res} . Then, the probability that the adversary A can distinguish between the set of runs in which the honest voters $(V_i)_{i \in \mathcal{I}_{hon}}$ vote according to $(v_i)_{i \in \mathcal{I}_{hon}}$ or to $(v'_i)_{i \in \mathcal{I}_{hon}}$ is negligible.*

Proof. We distinguish between the following two cases:

1. $(v_i)_{i \in \mathcal{I}_{hon}}$ and $(v'_i)_{i \in \mathcal{I}_{hon}}$ yield the result “no veto”.
2. $(v_i)_{i \in \mathcal{I}_{hon}}$ and $(v'_i)_{i \in \mathcal{I}_{hon}}$ yield the result “veto”.

In the first case, both $(v_i)_{i \in \mathcal{I}_{hon}}$ and $(v'_i)_{i \in \mathcal{I}_{hon}}$ consist of “no veto” choices only, hence $(v_i)_{i \in \mathcal{I}_{hon}} = (v'_i)_{i \in \mathcal{I}_{hon}}$. In particular, it is impossible to distinguish between runs in which the honest voters vote according to $(v_i)_{i \in \mathcal{I}_{hon}}$ or to $(v'_i)_{i \in \mathcal{I}_{hon}}$.

To prove indistinguishability in the second case, we use the following hybrid argument. To this end, we simulate the protocol as follows: if there exists at least one honest voter who chooses to veto, then *all* honest voters $(V_i)_{i \in \mathcal{I}_{hon}}$ veto. Under the assumption that $RLWE_{n,q,\chi}^{(m)}$ holds true, it follows that for any possible set of choices $(\tilde{v}_i)_{i \in \mathcal{I}_{hon}}$ which contains at least one “veto”, the simulated protocol is indistinguishable from the original veto protocol in which the honest voters vote according to $(\tilde{v}_i)_{i \in \mathcal{I}_{hon}}$. Due to the symmetry of this argument, we can conclude that no ppt adversary A can distinguish between runs in which the honest voters vote according to $(v_i)_{i \in \mathcal{I}_{hon}}$ or to $(v'_i)_{i \in \mathcal{I}_{hon}}$ if there exist $j, k \in \mathcal{I}_{hon}$ such that $v_j = \text{veto}$ and $v'_k = \text{veto}$.

5 Actively Secure Lattice-Based AV-Net

In this section, we describe how to extend the veto protocol from Section 4 such that it provides privacy and verifiable correctness in the presence of active adversaries.

Let us first explain why the protocol from Section 4 does neither protect privacy nor correctness if (some) voters do not follow their prescribed programs:

- *Privacy:* Assume that we have three voters V_1, V_2, V_3 , where V_1 and V_2 are honest, and V_3 is malicious and aims to actively break privacy of, say, voter V_1 . Now, V_3 waits until V_2 has published b_2 and then simply publishes $b_3 \leftarrow -b_2$. By this, we have that $y_1 = 0$. Hence, if V_1 does not veto, it follows

that $c_1 = e'_1$ is chosen according to χ , and that $c_1 = r_1$ is chosen uniformly at random otherwise. Therefore, the adversary (controlling V_3) knows that (with high probability) V_1 did not veto if $\|c_1\|_\infty < \beta$. This breaks V_1 's privacy.

- *Correctness*: Assume that we have two voters V_1, V_2 , where V_1 is honest and decides to veto, and V_2 is malicious and aims to actively cancel out V_1 's veto. Now, V_2 waits until V_1 has published c_1 and then simply publishes c_2 such that $\|c_1 + c_2\|_\infty < \frac{q}{4} - 2$. Therefore, the final result is “no veto” even though V_1 had chosen “veto”.

At a high level, what both attacks have in common is that the adversary can adaptively choose the corrupted voters outputs depending on the honest voters' ones. In order to eliminate this vulnerability, we employ a lattice-based commitment scheme as described in Section 5.1. We will then demonstrate that the resulting veto protocol in fact provides verifiable correctness (Section 5.2) and privacy (Section 5.3) against malicious adversaries.

5.1 Protocol description

We now explain how the passively secure veto protocol from Section 4 can be extended in order to defend against active adversary that aim to undermine privacy or verifiable correctness. More precisely, we need to ensure that voters choose their messages pairwise independently. To this end, we additionally employ an arbitrary lattice-based commitment scheme ($\text{KeyGen}_{\text{com}}, \text{Com}, \text{Open}$) which is (at least) computationally hiding and (at least) computationally binding under standard lattice hardness assumptions. More concretely, one could, for example, instantiate this generic commitment scheme with the highly efficient lattice-based commitment scheme by Baum et al. [2].

However, we need to be careful since commitment schemes like [2] are malleable. Even though there are generic compilers for transforming malleable commitment schemes into non-malleable ones (see, e.g., [8]), we are not aware of any existing work that analyzes such compilers in a quantum setting. Therefore, we will specify that voters open their commitments exactly in the reverse order according to which they published them. With this simple trick, we can still use malleable commitment schemes (see Section 7 for a discussion).

More precisely, we extend the veto protocol from Section 4 as follows. We refer to Appendix A for the notation related to the generic commitment scheme ($\text{KeyGen}_{\text{com}}, \text{Com}, \text{Open}$).

Parameters (extended). We denote by prm_{com} the joint public parameters of the commitment scheme (computed by running $\text{KeyGen}_{\text{com}}$).

Offline phase (extended). Each voter V_i , after having computed b_i , executes the following steps:

3. Compute $(\gamma_i, \rho_i) \leftarrow \text{Com}(\text{prm}_{\text{com}}, b_i)$.⁵

⁵ In other words, γ_i is the commitment to b_i using randomness ρ_i (see Appendix A).

4. Publish γ_i .
5. Wait until all γ_j were published ($j \in \{1, \dots, m\}$).
6. Set $\sigma \leftarrow$ order of published γ_j 's (according to their time stamps).
7. Wait until all (b_j, ρ_j) were published for $\sigma(j) > \sigma(i)$.
8. Publish (b_i, ρ_i) .
9. Wait until all (b_j, ρ_j) were published for $\sigma(j) < \sigma(i)$.
10. If $\text{Open}(\text{prm}_{\text{com}}, b_j, \gamma_j, \rho_j) = 0$ for some $j \neq i$, then abort.

Online phase (extended). Each voter V_i , after having computed c_i , executes the following steps:

3. Compute $(\gamma'_i, \rho'_i) \leftarrow \text{Com}(\text{prm}_{\text{com}}, c_i)$.
4. Publish γ'_i .
5. Wait until all γ'_j were published ($j \in \{1, \dots, m\}$).
6. Set $\sigma' \leftarrow$ order of published γ'_j 's (according to their time stamps).
7. Wait until all (c_j, ρ'_j) were published for $\sigma'(j) > \sigma'(i)$.
8. Publish (c_i, ρ'_i) .
9. Wait until all (c_j, ρ'_j) were published for $\sigma'(j) < \sigma'(i)$.
10. If $\text{Open}(\text{prm}_{\text{com}}, c_j, \gamma'_j, \rho'_j) = 0$ for some $j \neq i$, then abort.

5.2 Verifiable correctness

In this section, we show that the veto protocol defined in Section 5.1 is verifiably correct [7] even if an arbitrary adversary actively corrupts (a subset of) voters.

We note that we can restrict our attention to the case that an adversary aims to swap an honest “veto” into “no veto”. In fact, if an adversary (controlling at least one voter) wants the final result to be “veto”, then he can simply let the corrupted voter run her “veto” program.

Theorem 4 (Verifiable correctness). *Let P be the veto protocol defined in Section 5.1. Assume that the bulletin board B is honest. Assume that the commitment scheme is computationally binding and hiding. Then, we have that for all runs (of these instances) of P , the following implication holds true with overwhelming probability: If there exists an honest voter who chooses “veto”, then the final result is “veto” (or the protocol aborts prematurely).*

Proof. We assume w.l.o.g. that there exists one honest voter, namely, V_1 . This voter always chooses “veto”. Furthermore, we first restrict our attention to the case in which there exists one more voter, V_2 , which is controlled by an arbitrary ppt adversary A . Now, we distinguish between the following two sets of protocol runs:

1. Voter V_1 publishes γ'_1 before voter V_2 has published γ'_2 .
2. Voter V_2 publishes γ'_2 before voter V_1 has published γ'_1 .

In the first set of protocol runs, the probability that there is a final result and that this result is “no veto” equals to the probability that an arbitrary adversary A' can win the following game (run with challenger C):

1. C : choose $c_1 \xleftarrow{r} R_q$
2. C : compute $(\gamma_1, \rho_1) \leftarrow \text{Com}(\text{prm}_{\text{com}}, c_1)$
3. C : return γ_1
4. A' : return c_2
5. A' wins if and only if $\|c_1 + c_2\|_\infty < \frac{q}{4} - 2$

Since we assume that the commitment scheme is *computationally hiding* and since Lemma 4 holds true, any ppt A' can win this game only with at most negligible probability.

In the second set of protocol runs, the probability that there is a final result and that this result is “no veto” equals to the probability that an arbitrary adversary A' can win the following game (run with challenger C):

1. A' : return γ_2
2. C : choose $c_1 \xleftarrow{r} R_q$
3. C : return c_1
4. A' : return (c_2, ρ_2)
5. A' wins if and only if $\|c_1 + c_2\|_\infty < \frac{q}{4} - 2$ and $\text{Open}(\text{prm}_{\text{com}}, c_2, \gamma_2, \rho_2) = 1$.

Since we assume that the commitment scheme is *computationally binding* and since Lemma 4 holds true, any ppt A' can win this game only with at most negligible probability.

This proves that Theorem 4 holds true for one honest plus one dishonest voter. It is easy to see that the more general result, i.e., Theorem 4 with an arbitrary number of dishonest voters, effectively reduces to the two cases with one dishonest voter discussed above.

5.3 Privacy

In this section, we show that the veto protocol, as defined in Section 5.1, provides privacy in the presence of malicious adversaries.

Theorem 5 (Privacy). *Assume that $RLWE_{n,q,\chi}^{(m)}$ holds true. Assume that the commitment scheme is computationally binding and hiding. Let A be an arbitrary malicious ppt adversary which controls (at most) all but two voters $(V_i)_{i \in \mathcal{I}_{dis}}$. Let $(V_i)_{i \in \mathcal{I}_{hon}}$ denote the remaining (uncorrupted) voters. Let $(v_i)_{i \in \mathcal{I}_{hon}}$ and $(v'_i)_{i \in \mathcal{I}_{hon}}$ be two arbitrary vectors of choices that yield the same result res . Then, the probability that the adversary A can distinguish between the set of runs in which the honest voters $(V_i)_{i \in \mathcal{I}_{hon}}$ vote according to $(v_i)_{i \in \mathcal{I}_{hon}}$ or to $(v'_i)_{i \in \mathcal{I}_{hon}}$ is negligible.*

Proof. Using a similar argument as in the proof of Theorem 4, any adversary’s advantage of winning the privacy game (i.e., being able to distinguish between runs with $(v_i)_{i \in \mathcal{I}_{hon}}$ or with $(v'_i)_{i \in \mathcal{I}_{hon}}$) in the veto protocol from Section 5.1 is negligibly close to any adversary’s advantage of winning the privacy game in the following modification of this protocol. In fact, we modify the offline phase of protocol from Section 5.1 such that it is the same as the one of the passively

secure one (Section 4) but where first all dishonest voters publish their b_i 's and afterwards all honest ones. Now, even if all corrupted voters are (potentially) malicious, it follows from the privacy proof of the passively secure veto protocol (Theorem 3) that any adversary's advantage in winning the privacy game in this modified version is negligible under $\text{RLWE}_{n,q,\chi}^{(m)}$.

6 Experimental results

We have implemented the passively secure AV-net described in Section 4. Since the commitment scheme that is additionally required in the actively secure AV-net (Section 5) is generic and independent of the rest of the protocol, any efficient lattice-based commitment scheme can be chosen (e.g., [2]).

Our implementation uses C++ language and NTL library. We run 10,000 times experiments using the parameters (the same as in [19]) $n = 512, \sigma = 4.19, q = 120833$ on a computer with Intel Core i7-6500U CPU @ 2.50 GHz, running Cygwin version 3.1.5, g++ compiler version 9.3.0. Then we evaluate average runtime for discrete Gaussian sampling based on [21] (TimeDGS), polynomial multiplication (TimePoly), and vote tallying (TimeVeto) respectively. We show the experimental results with two decimal precision in Table 1.

Table 1: Runtime (millisecond) of our implementation.

m	TimeDGS	TimePoly	TimeVeto
3	0.42	0.89	0.24
10	2.44	8.44	6.94
15	5.60	23.49	16.51
20	8.69	39.44	24.35

The optimizer used was -O2. GCC basically performs almost all the supported optimizations that do not involve a space-speed tradeoff. This option is to benefit the compilation time and performance of the generated code. -O2 flags the compiler mainly to inline functions when able. -O3 adds some flags for loop unrolling and tree distribution and -Ofast disregards standards compliance and adds a couple extra flags like -ffast-math.

We just tested this code and it also works with -O3 as well as -Ofast, but at $m = 3$ and 10,000 runs, it does not appear to have any noticeable impact on the execution time of the code. We also tried several values for m and experimentally, no error showed up when $m = 100$ but errors start to show up when m is approximately 125.

7 Discussion

In this section, we elaborate on the properties of the AV-nets proposed and analyzed above.

Post-quantum anonymity. We have proven that the 2-round AV-net (Section 4.1) and the 4-round AV-net (Section 5.1) guarantee anonymity under the decisional RLWE assumption in the presence of arbitrary passive or active adversaries, respectively. The decisional RLWE assumption is a well-studied lattice-based hardness assumptions and commonly believed to be intractable even by quantum algorithms. Since anonymity of previous AV-nets [1, 13] relies on the DDH-assumption, our AV-nets are the first ones with *post-quantum* anonymity.

Observe that, both the two AV-nets proposed in this work as well as the previous one by Hao and Zielinski [13] have the following property: if there is a single voter who vetoes, then this voter knows that she is the only one who vetoed.

Robustness. It is obvious that if just a single voter does not participate in the online phase of our AV-net(s), then the complete protocol needs to restart again. Therefore, similarly to previous AV-nets [1, 13], our protocols have a low level of robustness, too. Typically, in order to increase robustness, protocols for secure computation employ threshold schemes: if at least t out of n parties participate, then the protocol terminates successfully. On the downside, however, threshold schemes lead to stronger trust assumptions for anonymity/privacy. In the case of (our) AV-nets, where we merely require that two voters are honest for anonymity, introducing a threshold structure would impair this mild trust assumption.

We note that in our actively secure protocol, opening the commitments in reverse order puts some burden on the underlying infrastructure, more precisely on the bulletin board. In fact, it is a non-trivial challenge in practice to guarantee verifiable time-stamps. One possible solution to this problem is to employ a distributed ledger technology (DLT).

Round complexity. Previous AV-nets [1, 13] require 2 rounds of interaction, both in the presence of passive and active adversaries. In contrast to that, our actively secure AV-net requires 4 rounds of interaction. The reason for this are the different techniques to make the voters' intrinsically homomorphic outputs *non-malleable*. While [1, 13] employ ZKPs for this purpose, it is not immediately clear how to efficiently do this in the lattice-based setting. Therefore, we decided to add two further rounds of interaction in which the voters first commit to their outputs before revealing them. Since there are a number of highly efficient lattice-based commitment schemes (see, e.g., [2]), we argue that our variant is a reasonable trade-off.

Alternative approaches. AV-nets can be regarded as specific instances of secure boardroom voting or, more generally, secure multi-party computation (MPC) protocols. We elaborate on this in what follows.

There are numerous efficient MPC protocols in the literature that could be used for securely evaluating veto functions, in particular with post-quantum privacy (see, e.g., [9]). Typically, employing such generic MPC protocols is advantageous for *complex* result functions. However, generic MPC protocols are less well-suited for the specific case of veto protocols, where the result function is simply Boolean OR.

In a boardroom voting protocol, the voters themselves tally the ballots, without having to rely on a trusted set of talliers or election authorities. Several such protocols have been proposed so far (see, e.g., [12, 14]). However, these protocols employ specific ZKPs, and therefore, as explained above, transforming them into a lattice-based setting undermines efficiency. Furthermore, we note that if we applied one of these boardroom voting protocols to evaluate the veto function, then the final result would reveal how many voters actually vetoed. In contrast to that, in an AV-net, the final result merely reveals whether or not at least one voter vetoed (without revealing the number of vetoing voters). Hence, AV-nets are *tally-hiding* [15] and thus provide an essentially perfect privacy level.

We note that existing verifiable post-quantum secure e-voting systems [4, 10] would not be (immediately) useful for our purposes as well. The reason is that they are neither tally-hiding nor designed for peer-to-peer elections.

8 Conclusion

We proposed the first AV-nets with post-quantum anonymity. The first variant of our protocol requires 2 rounds of interaction and is passively secure, whereas the second one requires 4 rounds of interaction and is actively secure. Anonymity of our AV-net reduces to the decisional ring learning with errors (RLWE) assumption.

Acknowledgements

We thank the anonymous reviewers for their constructive feedback. We also thank Peter B. Rønne for helpful remarks. Peter Y.A. Ryan and Johannes Müller acknowledge support from the Luxembourg National Research Fund (FNR) and the Research Council of Norway for the joint INTER project SURCVS (Number 11747298). Jintai Ding, Doug Emery, and Vonn Kee Wong would like to thank the US Air Force and NSF for partial support. Jintai Ding would also like to thank University of Luxembourg for partial support.

Bibliography

- [1] Samiran Bag, Muhammad Ajmal Azad, and Feng Hao. PriVeto: A Fully Private Two-Round Veto Protocol. *IET Information Security*, 13(4):311–320, 2019.
- [2] Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More Efficient Commitments from Structured Lattice Assumptions. In *SCN 2018, Proceedings*, volume 11035 of *LNCS*, pages 368–385. Springer, 2018.
- [3] David Bernhard, Véronique Cortier, David Galindo, Olivier Pereira, and Bogdan Warinschi. SoK: A Comprehensive Analysis of Game-Based Ballot Privacy Definitions. In *2015 IEEE S&P*, pages 499–516, 2015.
- [4] Xavier Boyen, Thomas Haines, and Johannes Müller. A Verifiable and Practical Lattice-Based Decryption Mix Net with External Auditing. In *ESORICS 2020*. To appear.
- [5] David Chaum. Security Without Identification: Transaction Systems to Make Big Brother Obsolete. *Commun. ACM*, 28(10):1030–1044, 1985.
- [6] David Chaum. The Dining Cryptographers Problem: Unconditional Sender and Recipient Untraceability. *J. Cryptology*, 1(1):65–75, 1988.
- [7] Véronique Cortier, David Galindo, Ralf Küsters, Johannes Müller, and Tomasz Truderung. SoK: Verifiability Notions for E-Voting Protocols. In *2016 IEEE S&P*, pages 779–798, 2016.
- [8] Giovanni Di Crescenzo, Yuval Ishai, and Rafail Ostrovsky. Non-Interactive and Non-Malleable Commitment. In *ACM STOC, 1998*, pages 141–150. ACM, 1998.
- [9] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty Computation from Somewhat Homomorphic Encryption. In *CRYPTO 2012. Proceedings*, volume 7417 of *LNCS*, pages 643–662. Springer, 2012.
- [10] Rafaël del Pino, Vadim Lyubashevsky, Gregory Neven, and Gregor Seiler. Practical Quantum-Safe Voting from Lattices. In *Proceedings of the 2017 ACM CCS*, pages 1565–1581, 2017.
- [11] Philippe Golle and Ari Juels. Dining Cryptographers Revisited. In *EUROCRYPT 2004, Proceedings*, volume 3027 of *LNCS*, pages 456–473. Springer, 2004.
- [12] Jens Groth. Efficient Maximal Privacy in Boardroom Voting and Anonymous Broadcast. In *FC 2004. Revised Papers*, volume 3110 of *LNCS*, pages 90–104. Springer, 2004.
- [13] Feng Hao and Piotr Zielinski. A 2-Round Anonymous Veto Protocol. In *Security Protocols, 14th International Workshop, 2006, Revised Selected Papers*, volume 5087 of *LNCS*, pages 202–211. Springer, 2006.
- [14] Aggelos Kiayias and Moti Yung. Self-tallying Elections and Perfect Ballot Secrecy. In *PKC 2002, Proceedings*, volume 2274 of *LNCS*, pages 141–158. Springer, 2002.

- [15] Ralf Küsters, Juliad Liedtke, Johannes Müller, Daniel Rausch, and Andreas Vogt. Ordinos: A Verifiable Tally-Hiding E-Voting System. In *IEEE EuroS&P 2020. To appear*, 2020.
- [16] Noah Stephens-Davidowitz. Discrete Gaussian sampling reduces to CVP and SVP. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1748-1764. Society for Industrial and Applied Mathematics, 2016.
- [17] Jintai Ding, Xiang Xie, Xiaodong Lin. A Simple Provably Secure Key Exchange Scheme Based on the Learning with Errors Problem. In *IACR Cryptology ePrint Archive*, Report 2012/688, 2012.
- [18] Jintai Ding, Tsuyoshi Takagi, Xinwei Gao, and Yuntao Wang. One Sample Ring-LWE with Rounding and Its Application to Key Exchange. In *Applied Cryptography and Network Security*, pages 323-343, Springer, 2019.
- [19] Jintai Ding, Tsuyoshi Takagi, Xinwei Gao, and Yuntao Wang. Ding Key Exchange. Technical report, National Institute of Standards and Technology, 2017. available at <https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Round-1-Submissions>.
- [20] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On Ideal Lattices and Learning with Errors over Rings. In *EUROCRYPT 2010. Proceedings*, volume 6110 of *LNCS*, pages 1-23. Springer, 2010.
- [21] Chris Peikert. An Efficient and Parallel Gaussian Sampler for Lattices. In *Advances in Cryptology CRYPTO 2010, 30th Annual Cryptology Conference*, pages 80-97, 2014.

A Commitment Schemes

A *commitment scheme* is a tuple of algorithms $(\text{KeyGen}_{\text{com}}, \text{Com}, \text{Open})$ where:

- $\text{KeyGen}_{\text{com}}$ is a ppt algorithm which takes 1^ℓ and outputs the public parameters prm_{com} , containing a definition of the *message space* $M_{\text{com}} = M_{\text{com}}^\ell$, the *commitment space* $C_{\text{com}} = C_{\text{com}}^\ell$, and the *opening space* $R = R^\ell$.
- Com is a ppt algorithm which takes $\text{prm}_{\text{com}}, m \in M_{\text{com}}$ and outputs values $c \in C_{\text{com}}$ and $r \in R$.
- Open is a deterministic polynomial-time algorithm which takes $\text{prm}_{\text{com}}, m \in M_{\text{com}}, c \in C_{\text{com}}, r \in R$ and outputs a bit $b \in \{0, 1\}$.