Constant-Round Group Key Exchange from the Ring-LWE Assumption

Daniel Apon[1], Dana Dachman-Soled[2], Huijing Gong[2(✉)], and Jonathan Katz[2]

[1] National Institute of Standards and Technology, Gaithersburg, USA
daniel.apon@nist.gov
[2] University of Maryland, College Park, USA
danadach@ece.umd.edu, {gong,jkatz}@cs.umd.edu

Abstract. Group key-exchange protocols allow a set of N parties to agree on a shared, secret key by communicating over a public network. A number of solutions to this problem have been proposed over the years, mostly based on variants of Diffie-Hellman (two-party) key exchange. To the best of our knowledge, however, there has been almost no work looking at candidate *post-quantum* group key-exchange protocols.

Here, we propose a constant-round protocol for unauthenticated group key exchange (i.e., with security against a passive eavesdropper) based on the hardness of the Ring-LWE problem. By applying the Katz-Yung compiler using any post-quantum signature scheme, we obtain a (scalable) protocol for *authenticated* group key exchange with post-quantum security. Our protocol is constructed by generalizing the Burmester-Desmedt protocol to the Ring-LWE setting, which requires addressing several technical challenges.

Keywords: Ring learning with errors · Post-quantum cryptography · Group key exchange

1 Introduction

Protocols for (authenticated) key exchange are among the most fundamental and widely used cryptographic primitives. They allow parties communicating over an insecure public network to establish a common secret key, called a *session key*, permitting the subsequent use of symmetric-key cryptography for encryption and authentication of sensitive data. They can be used to instantiate so-called "secure channels" upon which higher-level cryptographic protocols often depend.

Most work on key exchange, beginning with the classical paper of Diffie and Hellman, has focused on two-party key exchange. However, many works have also explored extensions to the *group* setting [1,2,5,6,8,9,11–17,21,22,24,25,29–31] in which N parties wish to agree on a common session key that they can each then use for encrypted/authenticated communication with the rest of the group.

The recent effort by NIST to evaluate and standardize one or more quantum-resistant public-key cryptosystems is entirely focused on digital signatures and

two-party key encapsulation/key exchange,[1] and there has been an extensive amount of research over the past decade focused on designing such schemes. In contrast, we are aware of almost no[2] work on *group* key-exchange protocols with post-quantum security beyond the observation that a post-quantum group key-exchange protocol can be constructed from any post-quantum two-party protocol by having a designated group manager run independent two-party protocols with the $N - 1$ other parties, and then send a session key of its choice to the other parties encrypted/authenticated using each of the resulting keys. Such a solution is often considered unacceptable since it is highly asymmetric, requires additional coordination, is not contributory, and puts a heavy load on a single party who becomes a central point of failure.

1.1   Our Contributions

In this work, we propose a constant-round group key-exchange protocol based on the hardness of the Ring-LWE problem [27], and hence with (plausible) post-quantum security. We focus on constructing an *unauthenticated* protocol—i.e., one secure against a passive eavesdropper—since known techniques such as the Katz-Yung compiler [24] can then be applied to obtain an *authenticated* protocol secure against an active attacker.

The starting point for our work is the two-round group key-exchange protocol by Burmester and Desmedt [15,16,24], which is based on the decisional Diffie-Hellman assumption. Assume a group $\mathbb{G}$ of prime order q and a generator $g \in \mathbb{G}$ are fixed and public. The Burmester-Desmedt protocol run by parties $P_0, \ldots, P_{N-1}$ then works as follows:

1. In the first round, each party P_i chooses uniform $r_i \in \mathbb{Z}_q$ and broadcasts $z_i = g^{r_i}$ to all other parties.
2. In the second round, each party P_i broadcasts $X_i = (z_{i+1}/z_{i-i})^{r_i}$ (where the parties' indices are taken modulo N).

Each party P_i can then compute its session key sk_i as

$$\mathsf{sk}_i = (z_{i-1})^{Nr_i} \cdot X_i^{N-1} \cdot X_{i+1}^{N-2} \cdots X_{i+N-2}.$$

One can check that all the keys are equal to the same value $g^{r_0 r_1 + \cdots + r_{N-1} r_0}$.

In attempting to adapt their protocol to the Ring-LWE setting, we could fix a ring R_q and a uniform element $a \in R_q$. Then:

1. In the first round, each party P_i chooses "small" secret value $s_i \in R_q$ and "small" noise term $e_i \in R_q$ (with the exact distribution being unimportant in the present discussion), and broadcasts $z_i = as_i + e_i$ to the other parties.

---

[1] Note that CPA-secure key encapsulation is equivalent to two-round key-exchange (with passive security).

[2] The protocol of Ding et al. [19] has no security proof; the work of Boneh et al. [10] shows a framework for constructing a group key-exchange protocol with plausible post-quantum security but without a concrete instantiation.

2. In the second round, each party P_i chooses a second "small" noise term $e_i' \in R_q$ and broadcasts $X_i = (z_{i+1} - z_{i-i}) \cdot s_i + e_i'$.

Each party can then compute a session key b_i as

$$b_i = N \cdot s_i \cdot z_{i-1} + (N-1) \cdot X_i + (N-2) \cdot X_{i+1} + \cdots + X_{i+N-2}.$$

The problem, of course, is that (due to the noise terms) these session keys computed by the parties will *not* be equal. They will, however, be "close" to each other if the $\{s_i, e_i, e_i'\}$ are all sufficiently small, so we can add an additional reconciliation step to ensure that all parties agree on a common key k.

This gives a protocol that is correct, but proving security (even for a passive eavesdropper) is more difficult than in the case of the Burmester-Desmedt protocol. Here we informally outline the main difficulties and how we address them. First, we note that trying to prove security by direct analogy to the proof of security for the Burmester-Desmedt protocol (cf. [24]) fails; in the latter case, it is possible to use the fact that, for example,

$$(z_2/z_0)^{r_1} = z_1^{r_2 - r_0},$$

whereas in our setting the analogous relation does not hold. In general, the natural proof strategy here is to switch all the $\{z_i\}$ values to uniform elements of R_q, and similarly to switch the $\{X_i\}$ values to uniform subject to the constraint that their sum is approximately 0 (i.e., subject to the constraint that $\sum_i X_i \approx 0$). Unfortunately this cannot be done by simply invoking the Ring-LWE assumption $O(N)$ times; in particular, the first time we try to invoke the assumption, say on the pair $(z_1 = as_1 + e_1, \; X_1 = (z_2 - z_0) \cdot s_1 + e_1')$, we need $z_2 - z_0$ to be uniform—which, in contrast to the analogous requirement in the Burmester-Desmedt protocol (for the value z_2/z_0), is not the case here. Thus, we must somehow break the circularity in the mutual dependence of the $\{z_i, X_i\}$ values.

Toward this end, let us look more carefully at the distribution of $\sum_i X_i$. We may write

$$\sum_i X_i = \sum_i (e_{i+1}s_i - e_{i-1}s_i) + \sum_i e_i'.$$

Consider now changing the way X_0 is chosen: that is, instead of choosing $X_0 = (z_1 - z_{N-1})s_0 + e_0'$ as in the protocol, we instead set $X_0 = -\sum_{i=1}^{N-1} X_i + e_0'$ (where e_0' is from the same distribution as before). Intuitively, as long as the standard deviation of e_0' is large enough, these two distributions of X_0 should be "close" (as they both satisfy $\sum_i X_i \approx 0$). This, in particular, means that we need the distribution of e_0' to be different from the distribution of the $\{e_i'\}_{i>0}$, as the standard deviation of the former needs to be larger than the latter.

We can indeed show that when we choose e_0' from an appropriate distribution then the Rényi divergence between the two distributions of X_0, above, is bounded by a polynomial. With this switch in the distribution of X_0, we have broken the circularity and can now use the Ring-LWE assumption to switch the distribution of z_0 to uniform, followed by the remaining $\{z_i, X_i\}$ values.

Unfortunately, bounded Rényi divergence does not imply statistical closeness. However, polynomially bounded Rényi divergence *does* imply that any event

occurring with negligible probability when X_0 is chosen according to the second distribution also occurs with negligible probability when X_0 is chosen according to the first distribution. For these reasons, we change our security goal from an "indistinguishability-based" one (namely, requiring that, given the transcript, the real session key is indistinguishable from uniform) to an "unpredictability-based" one (namely, given the transcript, it should be infeasible to compute the real session key). In the end, though, once the parties agree on an unpredictable value k they can hash it to obtain the final session key $\mathsf{sk} = \mathcal{H}(k)$; this final value sk will be indistinguishable from uniform if $\mathcal{H}$ is modeled as a random oracle.

2   Preliminaries

2.1   Notation

Let $\mathbb{Z}$ be the ring of integers, and let $[N] = \{0, 1, \ldots, N-1\}$. If χ is a probability distribution over some set S, then $x_0, x_1, \ldots, x_{\ell-1} \leftarrow \chi$ denotes independently sampling each x_i from distribution χ. We let $\mathrm{Supp}(\chi) = \{x : \chi(x) \neq 0\}$. Given an event E, we use $\overline{E}$ to denote its complement. Let $\chi(E)$ denote the probability that event E occurs under distribution χ. Given a polynomial p_i, let $(p_i)_j$ denote the jth coefficient of p_i. Let $\log(X)$ denote $\log_2(X)$, and $\exp(X)$ denote e^X.

2.2   Ring Learning with Errors

Informally, the (decisional) version of the Ring Learning with Errors (Ring-LWE) problem is: for some secret ring element s, distinguish many random "noisy ring products" with s from elements drawn uniform from the ring. More precisely, the Ring-LWE problem is parameterized by (R, q, χ, ℓ) as follows:

1. R is a ring, typically written as a polynomial quotient ring $R = \mathbb{Z}[X]/(f(X))$ for some irreducible polynomial $f(X)$ in the indeterminate X. In this paper, we restrict to the case of that $f(X) = X^n + 1$ where n is a power of 2. In later sections, we let R be parameterized by n.
2. q is a modulus defining the quotient ring $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$. We restrict to the case that q is prime and $q = 1 \mod 2n$.
3. $\chi = (\chi_s, \chi_e)$ is a pair of noise distributions over R_q (with χ_s the *secret key* distribution and χ_e the *error* distribution) that are concentrated on "short" elements, for an appropriate definition of "short" (e.g., the Euclidean distance metric on the integer-coefficients of the polynomials s or e drawn from R_q); and
4. ℓ is the number of samples provided to the adversary.

Formally, the Ring-LWE problem is to distinguish between ℓ samples independently drawn from one of two distributions. The first distribution is generated by fixing a random secret $s \leftarrow \chi_s$ then outputting

$$(a_i, b_i = s \cdot a_i + e_i) \in R_q \times R_q,$$

for $i \in [\ell]$, where each $a_i \in R_q$ is drawn uniformly at random and each $e_i \leftarrow \chi_e$ is drawn from the error distribution. For the second distribution, each sample $(a_i, b_i) \in R_q \times R_q$ is simply drawn uniformly at random.

Let A_{n,q,χ_s,χ_e} be the distribution that outputs the Ring-LWE sample $(a_i, b_i = s \cdot a_i + e_i)$ as above. We denote by $\mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_s,\chi_e,\ell}(\mathcal{B})$ the advantage of algorithm $\mathcal{B}$ in distinguishing distributions $A^{\ell}_{n,q,\chi_s,\chi_e}$ and $\mathcal{U}^{\ell}(R_q^2)$.

We define $\mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_s,\chi_e,\ell}(t)$ to be the maximum advantage of any adversary running in time t. Note that in later sections, we write as $\mathsf{Adv}_{n,q,\chi,\ell}$ when $\chi = \chi_s = \chi_e$ for simplicity.

The Ring-LWE Noise Distribution. The noise distribution χ (here we assume $\chi_s = \chi_e$, though this is not necessary) is usually a discrete Gaussian distribution on $R_q^{\vee}$ or in our case R_q (see [18] for details of the distinction, especially for concrete implementation purposes). Formally, in case of power of two cyclotomic rings, the discrete Gaussian distribution can be sampled by drawing each coefficient independently from the 1-dimensional discrete Gaussian distribution over $\mathbb{Z}$ with parameter σ, which is supported on $\{x \in \mathbb{Z} : -q/2 \leq x \leq q/2\}$ and has density function

$$D_{\mathbb{Z}_q,\sigma}(x) = \frac{e^{\frac{-\pi x^2}{\sigma^2}}}{\sum_{x=-\infty}^{\infty} e^{\frac{-\pi x^2}{\sigma^2}}}.$$

2.3 Rényi Divergence

The Rényi divergence (RD) is a measure of closeness of two probability distributions. For any two discrete probability distributions P and Q such that $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$, we define the Rényi divergence of order 2 as

$$\mathrm{RD}_2(P\|Q) = \sum_{x \in \mathrm{Supp}(P)} \frac{P(x)^2}{Q(x)}.$$

Rényi divergence has a probability preservation property that can be considered the multiplicative analogues of statistical distance.

Proposition 1. *Given discrete distributions P and Q with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$, let $E \in \mathrm{Supp}(Q)$ be an arbitrary event. We have*

$$Q(E) \geq P(E)^2/\mathrm{RD}_2(P\|Q).$$

This property implies that as long as $\mathrm{RD}_2(P\|Q)$ is bounded by $\mathrm{poly}(\lambda)$, any event E that occurs with negligible probability $Q(E)$ under distribution Q also occurs with negligible probability $P(E)$ under distribution P. We refer to [26,27] for the formal proof.

Theorem 2.1 ([7]). *Fix $m, q \in \mathbb{Z}$, a bound B, and the 1-dimensional discrete Gaussian distribution $D_{\mathbb{Z}_q, \sigma}$ with parameter σ such that $B < \sigma < q$. Moreover, let $e \in \mathbb{Z}$ be such that $|e| \leq B$. If $\sigma = \Omega(B\sqrt{m/\log \lambda})$, then*

$$\mathrm{RD}_2((e + D_{\mathbb{Z}_q, \sigma})^m || D_{\mathbb{Z}_q, \sigma}^m) \leq \exp(2\pi m(B/\sigma)^2) = \mathrm{poly}(\lambda),$$

where X^m denotes m independent samples from X.

2.4   Generic Key Reconciliation Mechanism

In this subsection, we define a generic, one round, two-party key reconciliation mechanism which allows both parties to derive the same key from an approximately agreed upon ring element. A key reconciliation mechanism KeyRec consists of two algorithms recMsg and recKey, parameterized by security parameter 1^λ as well as β_{Rec}. In this context, Alice and Bob hold "close" keys – b_A and b_B, respectively – and wish to generate a shared key k so that $k = k_A = k_B$. The abstract mechanism KeyRec is defined as follows:

1. Bob computes $(K, k_B) = \mathsf{recMsg}(b_B)$ and sends the reconciliation message K to Alice.
2. Once receiving K, Alice computes $k_A = \mathsf{recKey}(b_A, K) \in \{0, 1\}^\lambda$.

CORRECTNESS. Given $b_A, b_B \in R_q$, if each coefficient of $b_B - b_A$ is bounded by β_{Rec} – namely, $|b_B - b_A| \leq \beta_{\mathsf{Rec}}$ – then it is guaranteed that $k_A = k_B$.

SECURITY. A key reconciliation mechanism KeyRec is secure if the subsequent two distribution ensembles are computationally indistinguishable. (First, we describe a simple, helper distribution).

$\mathsf{Exe}_{\mathsf{KeyRec}}(\lambda)$: A draw from this helper distribution is performed by initiating the key reconciliation protocol among two honest parties and outputting (K, k_B); i.e. the reconciliation message K and (Bob's) key k_B of the protocol execution.

   We denote by $\mathsf{Adv}_{\mathsf{KeyRec}}(\mathcal{B})$ the advantage of adversary $\mathcal{B}$ distinguishing the distributions below.

$$\{(K, k_B) : b_B \leftarrow \mathcal{U}(R_q), (K, k_B) \leftarrow \mathsf{Exe}_{\mathsf{KeyRec}}(\lambda, b_B)\}_{\lambda \in \mathbb{N}},$$
$$\{(K, k') : b_B \leftarrow \mathcal{U}(R_q), (K, k_B) \leftarrow \mathsf{Exe}_{\mathsf{KeyRec}}(\lambda, b_B), k' \leftarrow U_\lambda\}_{\lambda \in \mathbb{N}},$$

where U_λ denotes the uniform distribution over λ bits.

   We define $\mathsf{Adv}_{\mathsf{KeyRec}}(t)$ to be the maximum advantage of any adversary running in time t.

Key Reconciliation Mechanisms from the Literature. The notion of key reconciliation was first introduced by Ding et al. [19]. in his work on two-party, lattice-based key exchange. It was later used in several important works on two-party key exchange, including [4,28,32].

   In the key reconciliation mechanisms of Peikert [28], Zhang et al. [32] and Alkim et al. [4], the initiating party sends a small amount of information about

its secret, b_B, to the other party. This information is enough to allow the two parties to agree upon the same key $k = k_A = k_B$, while revealing no information about k to an eavesdropper. When instantiating our GKE protocol with this type of key reconciliation (specifically, one of [4,28,32]), our final GKE protocol is "contributory," in the sense that all parties contribute entropy towards determining the final key.

Another method for the two parties to agree upon the same joint key $k = k_A = k_B$, given that they start with keys b_A, b_B that are "close," was first introduced in [3] (we refer to their technique as a key reconciliation mechanism, although it is technically not referred to as such in the literature). Here, the initiating party uses its private input to generate a Regev-style encryption of a random bit string k_B of its choice under secret key b_B. and then sends to the other party, who decrypts with its approximate secret key b_A to obtain k_A. Due to the inherent robustness to noise of Regev-style encryption, it is guaranteed that $k = k_A = k_B$ with all but negligible probability. Instantiating our GKE protocol with this type of key reconciliation (specifically, that in [3]) is also possible, but does not lead to the preferred "contributory GKE," since the initiating party's entropy completely determines the final group key.

3 Group Key Exchange Security Model

A group key-exchange protocol allows a session key to be established among $N > 2$ parties. Following prior work [12–14,23], we will use the term group key exchange (GKE) to denote a protocol secure against a *passive* (eavesdropping) adversary and will use the term authenticated group key exchange (GAKE) to denote a protocol secure against an *active* adversary, who controls all communication channels. Fortunately, the work of Katz and Yung [23] presents a compiler that takes any GKE protocol and transforms it into a GAKE protocol. The underlying tool required for this transform is any digital signature scheme which is strongly unforgeable under adaptive chosen message attack (EUF-CMA). We may thus focus our attention on achieving GKE in the remainder of this work.

In GKE, the adversary gets to see a single transcript generated by an execution of the GKE protocol. Given the transcript, the adversary must distinguish the real key from a fake key that is generated uniformly at random and independently of the transcript.

Formally, for security parameter $\lambda \in \mathbb{N}$, we define the following distribution:

$\mathsf{Execute}_\Pi^{\mathcal{O}_H}(\lambda)$: A draw from this distribution is performed by sampling a classical random oracle $\mathcal{H}$ from distribution $\mathcal{O}_H$, initiating the GKE protocol Π among N honest parties with security parameter λ relative to $\mathcal{H}$, and outputting $(\mathsf{trans}, \mathsf{sk})$—the transcript trans and key sk of the protocol execution.

Consider the following distributions:

$$\{(\mathsf{trans}, \mathsf{sk}) : (\mathsf{trans}, \mathsf{sk}) \leftarrow \mathsf{Execute}_\Pi^{\mathcal{O}_H}(\lambda)\}_{\lambda \in \mathbb{N}},$$
$$\{(\mathsf{trans}, \mathsf{sk}') : (\mathsf{trans}, \mathsf{sk}) \leftarrow \mathsf{Execute}_\Pi^{\mathcal{O}_H}(\lambda), \mathsf{sk}' \leftarrow U_\lambda\}_{\lambda \in \mathbb{N}},$$

where U_λ denotes the uniform distribution over λ bits. Let $\mathsf{Adv}^{\mathsf{GKE},\mathcal{O}_H}(\mathcal{A})$ denote the advantage of adversary $\mathcal{A}$, with classical access to the sampled oracle $\mathcal{H}$, distinguishing the distributions above.

To enable a concrete security analysis, we define $\mathsf{Adv}^{\mathsf{GKE},\mathcal{O}_H}(t, q_{\mathcal{O}_H})$ to be the maximum advantage of any adversary running in time t and making at most $q_{\mathcal{O}_H}$ queries to the random oracle. Security holds even if the adversary sees multiple executions by a hybrid argument.

In the next section we will define our GKE scheme and prove that it satisfies the notion of GKE.

4   A Group Key-Exchange Protocol

In this section, we present our group key exchange construction, GKE, which runs key reconciliation protocol KeyRec as a subroutine. Let KeyRec be parametrized by β_{Rec}. The protocol has two security parameters λ and ρ. λ is the computational security parameter, which is used in the security proof. ρ is the statistical security parameter, which is used in the correctness proof. σ_1, σ_2 are parameters of discrete Gaussian distributions. In this setting, N players $P_0, \ldots, P_{N-1}$ plan to generate a shared session key. The players' indices are taken modulo N.

The structure of the protocol is as follows: All parties agree on "close" keys $b_0 \approx \cdots \approx b_{N-1}$ after the second round. Player $N - 1$ then initiates a key reconciliation protocol to allow all users to agree on the same key $k = k_0 = \cdots = k_{N-1}$. Since we are only able to prove that k is difficult to compute for an eavesdropping adversary (but may not be indistinguishable from random), we hash k using random oracle $\mathcal{H}$ to get the final shared key sk.

Public parameter: $R_q = \mathbb{Z}_q[x]/(x^n + 1)$, $a \leftarrow \mathcal{U}(R_q)$.

Round 1: Each player P_i samples $s_i, e_i \leftarrow \chi_{\sigma_1}$ and broadcasts $z_i = as_i + e_i$.
Round 2: Player P_0 samples $e_0' \leftarrow \chi_{\sigma_2}$ and each of the other players P_i samples $e_i' \leftarrow \chi_{\sigma_1}$, broadcasts $X_i = (z_{i+1} - z_{i-1})s_i + e_i'$.
Round 3: Player P_{N-1} proceeds as follows:
 1. Samples $e_{N-1}'' \leftarrow \chi_{\sigma_1}$ and computes $b_{N-1} = z_{N-2}Ns_{N-1} + e_{N-1}'' + X_{N-1} \cdot (N-1) + X_0 \cdot (N-2) + \cdots + X_{N-3}$.
 2. Computes $(K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1})$ and broadcasts K_{N-1}.
 3. Obtains session key $\mathsf{sk}_{N-1} = \mathcal{H}(k_{N-1})$.
Key Computation: Each player P_i (except P_{N-1}) proceeds as follows:
 1. Computes $b_i = z_{i-1}Ns_i + X_i \cdot (N-1) + X_{i+1} \cdot (N-2) + \cdots + X_{i+N-2}$.
 2. Computes $k_i = \mathsf{recKey}(b_i, K_{N-1})$, and obtains session key $\mathsf{sk}_i = \mathcal{H}(k_i)$.

4.1   Correctness

The following claim states that each party derives the same session key sk_i, with all but negligible probability, as long as $\chi_{\sigma_1}, \chi_{\sigma_2}$ satisfy the constraint $(N^2 + 2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2} + 1)\sigma_1 + (N-2)\sigma_2 \leq \beta_{\mathsf{Rec}}$, where β_{Rec} is the parameter from the KeyRec protocol.

Theorem 4.1. *Given β_{Rec} as the parameter of* KeyRec *protocol, $N, n, \rho, \sigma_1, \sigma_2$ as parameters of GKE protocol Π, as long as $(N^2 + 2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2} + 1)\sigma_1 + (N - 2)\sigma_2 \leq \beta_{\mathsf{Rec}}$ is satisfied, if all players honestly execute the group key exchange protocol described above, then each player derives the same key as input of $\mathcal{H}$ with probability $1 - 2 \cdot 2^{-\rho}$.*

Proof. We refer to Sect. A of Appendix for the detailed proof. $\square$

5   Security Proof

The following theorem shows that protocol Π is a passively secure group key-exchange protocol. We remark that we prove security of the protocol for a classical attacker only; in particular, we allow the attacker only classical access to $\mathcal{H}$. We believe the protocol can be proven secure even against attackers that are allowed to make quantum queries to $\mathcal{H}$, but leave proving this to future work.

Theorem 5.1. *If the parameters in the group key exchange protocol Π satisfy the constraints $2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1 \leq \beta_{R\acute{e}nyi}$ and $\sigma_2 = \Omega(\beta_{R\acute{e}nyi}\sqrt{n/\log\lambda})$, and if $\mathcal{H}$ is modeled as a random oracle, then for any algorithm $\mathcal{A}$ running in time t, making at most q queries to the random oracle, we have:*

$$\mathsf{Adv}_{\Pi}^{\mathsf{GKE}, \mathcal{O}_H}(t, \mathsf{q}) \leq 2^{-\lambda+1}$$

$$+ \sqrt{\left(N \cdot \mathsf{Adv}_{n,q,\chi_{\sigma_1},3}^{\mathsf{RLWE}}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{\mathsf{q}}{2^\lambda}\right) \cdot \frac{\exp\left(2\pi n\left(\beta_{R\acute{e}nyi}/\sigma_2\right)^2\right)}{1 - 2^{-\lambda+1}}},$$

where $t_1 = t + \mathcal{O}(N) \cdot t_{\mathsf{ring}}, t_2 = t + \mathcal{O}(N) \cdot t_{\mathsf{ring}}$ and where t_{ring} is defined as the (maximum) time required to perform operations in R_q.

Proof. Consider the joint distribution of $(\mathsf{T}, \mathsf{sk})$, where $\mathsf{T} = (\{z_i\}, \{X_i\}, K_{k-1})$ is the transcript of an execution of the protocol Π, and k is the final shared session key. The distribution of $(\mathsf{T}, \mathsf{sk})$ is denoted as Real. Proceeding via a sequence of experiments, we will show that under the Ring-LWE assumption, if an efficient adversary queries the random oracle on input k_{N-1} in the Ideal experiment (to be formally defined) with at most negligible probability, then it also queries the random oracle on input k_{N-1} in the Real experiment with at most negligible probability.

Furthermore, in Ideal, the input k_{N-1} to the random oracle is uniform random, which means that the adversary has $\mathsf{negl}(\lambda)$ probability of guessing k_{N-1} in Ideal when $\mathsf{q} = \mathrm{poly}(\lambda)$. Finally, we argue that the above is sufficient to prove the GKE security of the scheme, because in the random oracle model, the output of the random oracle on k_{N-1} – i.e. the agreed upon key – looks uniformly random to an adversary who does not query k_{N-1}. We now proceed with the formal proof.

Let Query be the event that k_{N-1} is among the adversary $\mathcal{A}$'s random oracle queries and denote by $\mathrm{Pr}_i[\mathsf{Query}]$ the probability that event Query happens in *Experiment i*. Note that we let $e_0' = \hat{e}_0$ in order to distinguish this from the other e_i''s sampled from a different distribution.

Experiment 0. This is the original experiment. In this experiment, the distribution of $(\mathsf{T}, \mathsf{sk})$ is as follows, denoted Real:

$$
\mathsf{Real} := \left\{
\begin{array}{l}
a \leftarrow R_q; s_0, s_1, \ldots, s_{N-1}, e_0, e_1, \ldots, e_{N-1} \leftarrow \chi; \\
z_0 = as_0 + e_0, z_1 = as_1 + e_1, \ldots, z_{N-1} = as_{N-1} + e_{N-1}; \\
e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; \hat{e}_0 \leftarrow \chi_{\sigma_2}; \\
X_0 = (z_1 - z_{N-1})s_0 + \hat{e}_0, X_1 = (z_2 - z_0)s_1 + e'_1, \ldots, \\
X_{N-1} = (z_0 - z_{N-2})s_{N-1} + e'_{N-1}; e''_{N-1} \leftarrow \chi_{\sigma_1}; \qquad : (\mathsf{T}, \mathsf{sk}) \\
b_{N-1} = z_{N-2}Ns_{N-1} + e''_{N-1} + X_{N-1} \cdot (N-1) + \\
\quad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\
(K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\
\mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}).
\end{array}
\right\}.
$$

Since $\Pr[\mathcal{A}\text{ succeeds}] = \frac{1}{2} + \mathsf{Adv}_\Pi^{\mathsf{GKE}, \mathcal{O}_H}(t, \mathsf{q}) = \Pr_0[\mathsf{Query}] \cdot 1 + \Pr_0(\overline{\mathsf{Query}}) \cdot \frac{1}{2}$, we have

$$
\mathsf{Adv}_\Pi^{\mathsf{GKE}, \mathcal{O}_H}(t, \mathsf{q}) \leq \Pr_0[\mathsf{Query}]. \tag{1}
$$

In the remainder of the proof, we focus on bounding $\Pr_0[\mathsf{Query}]$.

Experiment 1. In this experiment, X_0 is replaced by $X'_0 = -\sum_{i=1}^{N-1} X_i + \hat{e}_0$. The remainder of the experiment is exactly the same as *Experiment 0*. The corresponding distribution of $(\mathsf{T}, \mathsf{sk})$ is as follows, denoted Dist_1:

$$
\mathsf{Dist}_1 := \left\{
\begin{array}{l}
a \leftarrow \mathcal{U}(R_q); s_0, s_1, \ldots, s_{N-1}, e_0, e_1, \ldots, e_{N-1} \leftarrow \chi_{\sigma_1}; \\
z_0 = as_0 + e_0, z_1 = as_1 + e_1, \ldots, z_{N-1} = as_{N-1} + e_{N-1}; \\
e'_0, e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; \hat{e}_0 \leftarrow \chi_{\sigma_2} \\
X'_0 = -\sum_{i=1}^{N-1} X_i + \hat{e}_0, X_1 = (z_2 - z_0)s_1 + e'_1, \ldots, \\
X_{N-1} = (z_0 - z_{N-2})s_{N-1} + e'_{N-1}; e''_{N-1} \leftarrow \chi_{\sigma_1}; \qquad : (\mathsf{T}, \mathsf{sk}) \\
b_{N-1} = z_{N-2}Ns_{N-1} + e''_{N-1} + X_{N-1} \cdot (N-1) + \\
\quad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\
(K_{N-1}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\
\mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, K_{N-1}).
\end{array}
\right\}.
$$

Claim. Given $a \leftarrow \mathcal{U}(R_q)$, $s_0, s_1, \ldots, s_{N-1}, e_0, e_1, \ldots, e_{N-1}, e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}$, $\hat{e}_0 \leftarrow \chi_{\sigma_2}$, $X_0 = (z_1 - z_{N-1})s_0 + \hat{e}_0$, $X'_0 = -\sum_{i=1}^{N-1} X_i + \hat{e}_0$, where $R_q, \chi_{\sigma_1}, \chi_{\sigma_2}$, $z_1, z_{N-1}, X_1, \ldots, X_{N-1}$ are defined as above, and the constraint $2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1 \leq \beta_{\mathsf{Rényi}}$ is satisfied, we have

$$
\Pr_0[\mathsf{Query}] \leq \sqrt{\Pr_1[\mathsf{Query}] \cdot \frac{\exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2)}{1 - 2^{-\lambda+1}}} + 2^{-\lambda+1}. \tag{2}
$$

Proof. Let $\mathsf{Error} = \sum_{i=0}^{N-1}(s_i e_{i+1} + s_i e_{i-1}) + \sum_{i=1}^{N-1} e_i'$. We begin by showing that the absolute value of each coefficient of Error is bounded by $\beta_{\mathsf{R\acute{e}nyi}}$ with all but negligible probability. Then by adding a "bigger" error $\hat{e}_0 \leftarrow \chi_{\sigma_2}$, the small difference between distributions $\mathsf{Error} + \chi_{\sigma_2}$ (corresponding to Experiment 0) and χ_{σ_2} (corresponding to Experiment 1) can be "washed" away by applying Theorem 2.1.

For all coefficient indices j, note that $|\mathsf{Error}_j| = |(\sum_{i=0}^{N-1}(s_i e_{i+1} + s_i e_{i-1}) + \sum_{i=1}^{N-1} e_i')_j|$. Let bound_λ denote the event that for all i and all coordinate indices j, $|(s_i)_j| \leq c\sigma_1$, $|(e_i)_j| \leq c\sigma_1$, $|(e_i')_j| \leq c\sigma_1$, $|(e_{N-1}'')_j| \leq c\sigma_1$, and $|(\hat{e}_0)_j| \leq c\sigma_2$, where $c = \sqrt{\frac{2\lambda}{\pi \log e}}$. By replacing ρ with λ in Lemmas A.1 and A.2 and by a union bound, we have – conditioned on bound_λ – that $|\mathsf{Error}_j| \leq 2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1$ for all j, with probability at least $1 - 2N \cdot 2n2^{-2\lambda}$. Since, under the assumption that $4Nn \leq 2^\lambda$, we have that $\Pr[\mathsf{bound}_\lambda] \geq 1 - 2^{-\lambda}$, we conclude that

$$\Pr[|\mathsf{Error}_j| \leq \beta_{\mathsf{R\acute{e}nyi}}, \forall j] \geq 1 - 2^{-\lambda+1}. \tag{3}$$

For a fixed $\mathsf{Error} \in R_q$, we denote by D_1 the distribution of $\mathsf{Error} + \chi_{\sigma_2}$ and note that D_1, χ_{σ_2} are n-dimension distributions.

Since $\sigma_2 = \Omega(\beta_{\mathsf{R\acute{e}nyi}}\sqrt{n/\log \lambda})$, assuming that for all j, $|\mathsf{Error}_j| \leq \beta_{\mathsf{R\acute{e}nyi}}$, by Theorem 2.1, we have

$$\mathrm{RD}_2(D_1 \| \chi_{\sigma_2}) \leq \exp(2\pi n(\beta_{\mathsf{R\acute{e}nyi}}/\sigma_2)^2) = \mathrm{poly}(\lambda). \tag{4}$$

Then it is straightforward to verify that the distribution of X_0 in *Experiment 0* is

$$\left(as_1 s_0 - as_{N-1} s_0 - \sum_{i=0}^{N-1}(e_{i+1} s_i + e_{i-1} s_i) - \sum_{i=1}^{N-1} e_i'\right) + D_1,$$

and the distribution of X_0' in *Experiment 1* is

$$\left(as_1 s_0 - as_{N-1} s_0 - \sum_{i=0}^{N-1}(e_{i+1} s_i + e_{i-1} s_i) - \sum_{i=1}^{N-1} e_i'\right) + \chi_{\sigma_2}.$$

In addition, the remaining part of Dist_1 is identical to Real. Therefore we may view Real in *Experiment 0* as a function of a random variable sampled from D_1 and take Dist_1 in *Experiment 1* as a function of a random variable sampled from χ_{σ_2}.

Recall that Query is the event that k_{N-1} is contained in the set of random oracle queries issued by adversary $\mathcal{A}$. We denote by Xbound the event that $|\mathsf{Error}_j| \leq \beta_{\mathsf{R\acute{e}nyi}}, \forall j$. Note that computation of Error_j is available in both *Experiment 0* and *Experiment 1*. We denote by $\Pr_0[\mathsf{Xbound}]$ (resp. $\Pr_1[\mathsf{Xbound}]$) the probability that event Xbound occurs in *Experiment 0* (resp. *Experiment 1*) and define $\Pr_0[\overline{\mathsf{Xbound}}], \Pr_1[\overline{\mathsf{Xbound}}]$ analogously. Let Real' (resp. Dist_1') denote the random variable Real (resp. Dist_1), conditioned on the event Xbound. Therefore, we have

$$\Pr_0[\mathsf{Query}] = \Pr_0[\mathsf{Query}|\mathsf{Xbound}] \cdot \Pr_0[\mathsf{Xbound}] + \Pr_0[\mathsf{Query}|\overline{\mathsf{Xbound}}] \cdot \Pr_0[\overline{\mathsf{Xbound}}]$$

$$\leq \Pr_0[\mathsf{Query}|\mathsf{Xbound}] + \Pr_0[\overline{\mathsf{Xbound}}]$$

$$\leq \Pr_0[\mathsf{Query}|\mathsf{Xbound}] + 2^{-\lambda+1}$$

$$\leq \sqrt{\Pr_1[\mathsf{Query}|\mathsf{Xbound}] \cdot \mathrm{RD}_2(\mathsf{Real}'||\mathsf{Dist}'_1)} + 2^{-\lambda+1}$$

$$\leq \sqrt{\Pr_1[\mathsf{Query}|\mathsf{Xbound}] \cdot \mathrm{RD}_2(D_1||\chi_{\sigma_2})} + 2^{-\lambda+1}$$

$$\leq \sqrt{\Pr_1[\mathsf{Query}|\mathsf{Xbound}] \cdot \exp(2\pi n(\beta_{\mathsf{R\acute{e}nyi}}/\sigma_2)^2)} + 2^{-\lambda+1}$$

$$\leq \sqrt{\Pr_1\left[\mathsf{Query} \cdot \frac{\exp(2\pi n(\beta_{\mathsf{R\acute{e}nyi}}/\sigma_2)^2)}{\Pr_1[\mathsf{Xbound}]}\right]} + 2^{-\lambda+1}$$

$$\leq \sqrt{\Pr_1[\mathsf{Query}] \cdot \frac{\exp(2\pi n(\beta_{\mathsf{R\acute{e}nyi}}/\sigma_2)^2)}{1 - 2^{-\lambda+1}}} + 2^{-\lambda+1},$$

where the second and last inequalities follow from (3), the third inequality follows from Proposition 1 and the fifth inequality follows from (4). □

Due to page restriction, we defer the proof of showing

$$\Pr_1[\mathsf{Query}] \leq \left(N \cdot \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{q}{2^\lambda}\right),$$

to the full version. □

5.1 Parameter Constraints

Beyond the parameter settings recommended for instantiating Ring-LWE with security parameter λ, parameters $N, n, \sigma_1, \sigma_2, \lambda, \rho$ of the protocol above are also required to satisfy the following inequalities:

$$(N^2 + 2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2} + 1)\sigma_1 + (N - 2)\sigma_2 \leq \beta_{\mathsf{Rec}} \quad \text{(Correctness)} \quad (5)$$

$$2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N - 1)\sigma_1 \leq \beta_{\mathsf{R\acute{e}nyi}} \quad \text{(Security)} \quad (6)$$

$$\sigma_2 = \Omega(\beta_{\mathsf{R\acute{e}nyi}}\sqrt{n/\log\lambda}) \quad \text{(Security)} \quad (7)$$

We comment that once the ring, the noise distributions, and the security parameters λ, ρ are fixed, the maximum number of parties is fixed.

A   Correctness of the Group Key-Exchange Protocol

Theorem 4.1. *Given β_{Rec} as parameter of KeyRec protocol, $N, n, \rho, \sigma_1, \sigma_2$ as parameters of GKE protocol Π, $(N^2+2N)\cdot\sqrt{n}\rho^{3/2}\sigma_1^2+(\frac{N^2}{2}+1)\sigma_1+(N-2)\sigma_2 \leq \beta_{\mathsf{Rec}}$ is satisfied, if all players honestly execute the group key exchange protocol as described above, then each player derive the same key as input of $\mathcal{H}$ with probability $1 - 2 \cdot 2^{-\rho}$.*

Proof. Given $s_i, e_i, e_i', e_{N-1}'' \leftarrow \chi_{\sigma_1}, \hat{e}_0 \leftarrow \chi_{\sigma_2}$ for all i as specified in protocol Π, we begin by introducing the following lemmas to analyze probabilities that each coordinate of $s_i, e_i, e_i', e_{N-1}'', \hat{e}_0$ are "short" for all i, and conditioned on the first event, $s_i e_i$ are "short".

Lemma A.1. *Given $s_i, e_i, e_i', e_{N-1}'', \hat{e}_0$ for all i as defined above, let bound denote the event that for all i and all coordinate indices j, $|(s_i)_j| \leq c\sigma_1$, $|(e_i)_j| \leq c\sigma_1$, $|(e_i')_j| \leq c\sigma_1$, $|(e_{N-1}'')_j| \leq c\sigma_1$, and $|(\hat{e}_0)_j| \leq c\sigma_2$, where $c = \sqrt{\frac{2\rho}{\pi \log e}}$, we have $\Pr[\mathsf{bound}] \geq 1 - 2^{-\rho}$.*

Proof. Using the fact that complementary error function $\mathrm{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt \leq e^{-x^2}$, we obtain

$$\Pr[|v| \geq c\sigma + 1; v \leftarrow D_{\mathbb{Z}_q,\sigma}] \leq 2 \sum_{x=\lfloor c\sigma+1 \rfloor}^\infty D_{\mathbb{Z}_q,\sigma}(x) \leq \frac{2}{\sigma} \int_{c\sigma}^\infty e^{-\frac{\pi x^2}{\sigma^2}} dx$$

$$= \frac{2}{\sqrt{\pi}} \int_{\frac{\sqrt{\pi}}{\sigma}(c\sigma)}^\infty e^{-t^2} dt \leq e^{-c^2\pi}.$$

Note that there are $3nN$ number of coordinates sampled from distribution $D_{\mathbb{Z}_q,\sigma_1}$, and n number of coordinates sampled from distribution $D_{\mathbb{Z}_q,\sigma_2}$ in total. Assume $3nN + n \leq e^{c^2\pi/2}$, since all the coordinates are sampled independently, we bound $\Pr[\mathsf{bound}]$ as follow:

$$\Pr[\mathsf{bound}] = \left(1 - \Pr[|v| \geq c\sigma_1 + 1; v \leftarrow D_{\mathbb{Z}_q,\sigma_1}]\right)^{3nN}$$
$$\cdot \left(1 - \Pr[|\hat{e}_0| \geq c\sigma_2 + 1; \hat{e}_0 \leftarrow D_{\mathbb{Z}_q,\sigma_2}]\right)^n$$
$$\geq 1 - (3nN + n)e^{-c^2\pi} \geq 1 - e^{-c^2\pi/2} \geq 1 - 2^{-\rho}.$$

The last inequality follows as $c = \sqrt{\frac{2\rho}{\pi \log e}}$. $\square$

Lemma A.2. *Given $s_i, e_i, e_i', e_{N-1}'', \hat{e}_0$ for all i as defined above, and bound as defined in Lemma A.1, let $\mathsf{product}_{s_i,e_j}$ denote the event that, for all coefficient indices v, $|(s_i e_j)_v| \leq \sqrt{n}\rho^{3/2}\sigma_1^2$. we have*

$$\Pr[\mathsf{product}_{s_i,e_j}|\mathsf{bound}] \geq 1 - 2n \cdot 2^{-2\rho}.$$

Proof. For $t \in \{0, \ldots, n-1\}$, Let $(s_i)_t$ denote the t^{th} coefficient of $s_i \in R_q$, namely, $s_i = \sum_{t=0}^{n-1} (s_i)_t X^i$. $(e_j)_t$ is defined analogously. Since we have $X^n + 1$ as modulo of R, it is easy to see that $(s_i e_j)_v = c_v X^v$, where $c_v = \sum_{u=0}^{n-1} (s_i)_u (e_j)^*_{v-u}$, and $(e_j)^*_{v-u} = (e_j)_{v-u}$ if $v - u \geq 0$, $(e_j)^*_{v-u} = -(e_j)_{v-u+n}$, otherwise. Thus, conditioned on $|(s_i)_t| \leq c\sigma_1$ and $|(e_j)_t| \leq c\sigma_1$ (for all i, j, t) where $c = \sqrt{\frac{2\rho}{\pi \log e}}$, by Hoeffding's Inequality [20], we derive

$$\Pr[|(s_i e_j)_v| \geq \delta] = \Pr\left[\left|\sum_{u=0}^{n-1}(s_i)_u(e_j)^*_{v-u}\right| \geq \delta\right] \leq 2\exp\left(\frac{-2\delta^2}{n(2c^2\sigma_1^2)^2}\right),$$

as each product $(s_i)_u (e_j)^*_{v-u}$ in the sum is an independent random variable with mean 0 in the range $[-c^2\sigma_1^2, c^2\sigma_1^2]$. By setting $\delta = \sqrt{n}\rho^{3/2}\sigma_1^2$, we obtain

$$\Pr[|(s_u e_v)_i| \geq \sqrt{n}\rho^{3/2}\sigma_1^2] \leq 2^{-2\rho+1}. \tag{8}$$

Finally, by Union Bound,

$$\Pr[\mathsf{product}_{s_i, e_j}|\mathsf{bound}] = \Pr[|(s_i e_j)_v| \leq \sqrt{n}\rho^{3/2}\sigma_1^2, \forall v] \geq 1 - 2n \cdot 2^{-2\rho}. \tag{9}$$

$\square$

Now we begin analyzing the chance that not all parties agree on the same final key. The correctness of KeyRec guarantees that this group key exchange protocol has agreed session key among all parties $\forall i, k_i = k_{N-1}$, if $\forall j$, the j^{th} coefficient of $|b_{N-1} - b_i| \leq \beta_{\mathsf{Rec}}$.

For better illustration, we first write $X_0, \ldots, X_{N-1}$ in form of linear system as follows. $\boldsymbol{X} = [X_0 \ X_1 \ X_2 \ \cdots \ X_{N-1}]^T$

$$= \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & \ldots & 0 & -1 \\ -1 & 1 & 0 & 0 & \ldots & 0 & 0 \\ 0 & -1 & 1 & 0 & \ldots & 0 & 0 \\ 0 & 0 & -1 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & & \ddots & & \vdots \\ 0 & 0 & 0 & 0 & \ldots & -1 & 1 \end{bmatrix}}_{\boldsymbol{M}} \underbrace{\begin{bmatrix} as_0s_1 \\ as_1s_2 \\ as_2s_3 \\ as_3s_4 \\ \vdots \\ as_{N-2}s_{N-1} \\ as_{N-1}s_0 \end{bmatrix}}_{\boldsymbol{S}} + \underbrace{\begin{bmatrix} s_0e_1 - s_0e_{N-1} + e_0' \\ s_1e_2 - s_1e_0 + e_1' \\ s_2e_3 - s_2e_1 + e_2' \\ s_3e_4 - s_3e_2 + e_3' \\ \vdots \\ s_{N-2}e_{N-3} - s_{N-2}e_{N-3} + e_{N-2}' \\ s_{N-1}e_0 - s_{N-1}e_{N-2} + e_{N-1}' \end{bmatrix}}_{\boldsymbol{E}}. \tag{10}$$

We denote the matrices above by $\boldsymbol{M}, \boldsymbol{S}, \boldsymbol{E}$ from left to right and have the linear system as $\boldsymbol{X} = \boldsymbol{MS} + \boldsymbol{E}$. By setting $\boldsymbol{B}_i = [i-1 \ i-2 \ \cdots \ 0 \ N-1 \ N-2 \ \cdots \ i]$ as a N-dimensional vector, we can then write b_i as $\boldsymbol{B}_i \cdot \boldsymbol{X} + N(as_i s_{i-1} + s_i e_{i-1}) = \boldsymbol{B}_i \boldsymbol{MS} + \boldsymbol{B}_i \boldsymbol{E} + N(as_i s_{i-1} + s_i e_{i-1})$, for $i \neq N-1$ and write b_{N-1} as $\boldsymbol{B}_{N-1}\boldsymbol{MS} + \boldsymbol{B}_{N-1}\boldsymbol{E} + N(as_{N-1}s_{N-2} + s_{N-1}e_{N-2}) + e_{N-1}''$. It is straightforward to see that, entries of $\boldsymbol{MS}$ and $Nas_i s_{i-1}$ are eliminated through the process of computing $b_{N-1} - b_i$. Thus we get

$$b_{N-1} - b_i = (\boldsymbol{B}_{N-1} - \boldsymbol{B}_i)\,\boldsymbol{E} + N(s_{N-1}e_{N-2} - s_i e_{i-1}) + e''_{N-1}$$

$$= (N - i - 1) \cdot \left(\sum_{\substack{j \in \mathbb{Z} \cap [0, i-1] \\ \text{and } j = N-1}} s_j e_{j+1} - s_j e_{j-1} + e'_j \right) + e''_{N-1}$$

$$+ (-i - 1) \left(\sum_{j=i}^{N-2} s_j e_{j+1} - s_j e_{j-1} + e'_j \right) + N(s_{N-1}e_{N-2} - s_i e_{i-1})$$

Observe that for an arbitrary $i \in [N]$, there are at most $(N^2 + 2N)$ terms in form of $s_u e_v$, at most $N^2/2$ terms in form of e'_w where $e'_w \leftarrow \chi_{\sigma_1}$, at most $N-2$ terms of e'_0, where $e'_0 \leftarrow \chi_{\sigma_2}$, and one term in form of e''_{N-1} in any coordinate of the sum above. Let $\mathsf{product}_{\mathsf{ALL}}$ denote the event that for all the terms in form of $s_u e_v$ observed above, each coefficient of such term is bounded by $\sqrt{n}\rho^{3/2}\sigma_1^2$. By Union Bound and by assuming $2n(N^2 + 2N) \leq 2^\rho$, it is straightforward to see $\Pr[\overline{\mathsf{product}_{\mathsf{ALL}}}|\mathsf{bound}] \leq (N^2 + 2N) \cdot 2n2^{-2\rho} \leq 2^{-\rho}$.

Let bad be the event that not all parties agree on the same final key. Given the constraint $(N^2 + 2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2} + 1)\sigma_1 + (N-2)\sigma_2 \leq \beta_{\mathsf{Rec}}$ satisfied, we have

$$\Pr[\mathsf{bad}] = \Pr[\mathsf{bad}|\mathsf{bound}] \cdot \Pr[\mathsf{bound}] + \Pr[\mathsf{bad}|\overline{\mathsf{bound}}] \cdot \Pr[\overline{\mathsf{bound}}] \tag{11}$$

$$\leq \Pr[\overline{\mathsf{product}_{\mathsf{ALL}}}] \cdot 1 + 1 \cdot \Pr[\overline{\mathsf{bound}}] \leq 2 \cdot 2^{-\rho}, \tag{12}$$

which completes the proof. □

References

1. Abdalla, M., Bresson, E., Chevassut, O., Pointcheval, D.: Password-based group key exchange in a constant number of rounds. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.) PKC 2006. LNCS, vol. 3958, pp. 427–442. Springer, Heidelberg (2006). https://doi.org/10.1007/11745853_28

2. Abdalla, M., Pointcheval, D.: A scalable password-based group key exchange protocol in the standard model. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 332–347. Springer, Heidelberg (2006). https://doi.org/10.1007/11935230_22

3. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: NewHope without reconciliation. Cryptology ePrint Archive, Report 2016/1157 (2016). http://eprint.iacr.org/2016/1157

4. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange—a new hope. In: 25th USENIX Security Symposium (USENIX Security 2016), pp. 327–343. USENIX Association, Austin (2016)

5. Becker, K., Wille, U.: Communication complexity of group key distribution. In: Proceedings of the 5th ACM Conference on Computer and Communications Security, CCS 1998, pp. 1–6. ACM, New York (1998)

6. Bellare, M., Rogaway, P.: Provably secure session key distribution: the three party case. In: 27th Annual ACM Symposium on Theory of Computing, Las Vegas, NV, USA, 29 May–1 June, pp. 57–66. ACM Press (1995)

7. Bogdanov, A., Guo, S., Masny, D., Richelson, S., Rosen, A.: On the hardness of learning with rounding over small modulus. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016. LNCS, vol. 9562, pp. 209–224. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49096-9_9

8. Bohli, J.-M., Vasco, M.I.G., Steinwandt, R.: Password-authenticated constant-round group key establishment with a common reference string. Cryptology ePrint Archive, Report 2006/214 (2006). http://eprint.iacr.org/2006/214

9. Bohli, J.-M., Vasco, M.I.G., Steinwandt, R.: Secure group key establishment revisited. Int. J. Inf. Secur. **6**(4), 243–254 (2007)

10. Boneh, D., et al.: Multiparty non-interactive key exchange and more from isogenies on elliptic curves. arXiv preprint arXiv:1807.03038 (2018)

11. Bresson, E., Catalano, D.: Constant round authenticated group key agreement via distributed computation. In: Bao, F., Deng, R., Zhou, J. (eds.) PKC 2004. LNCS, vol. 2947, pp. 115–129. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24632-9_9

12. Bresson, E., Chevassut, O., Pointcheval, D.: Provably authenticated group Diffie-Hellman key exchange—the dynamic case. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 290–309. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45682-1_18

13. Bresson, E., Chevassut, O., Pointcheval, D.: Dynamic group Diffie-Hellman key exchange under standard assumptions. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 321–336. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_21

14. Bresson, E., Chevassut, O., Pointcheval, D., Quisquater, J.-J.: Provably authenticated group Diffie-Hellman key exchange. In: 8th Conference on Computer and Communications Security, ACM CCS 2001, Philadelphia, PA, USA, 5–8 November, pp. 255–264. ACM Press (2001)

15. Burmester, M., Desmedt, Y.: A secure and efficient conference key distribution system. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 275–286. Springer, Heidelberg (1995). https://doi.org/10.1007/BFb0053443

16. Burmester, M., Desmedt, Y.: A secure and scalable group key exchange system. Inf. Process. Lett. **94**(3), 137–143 (2005)

17. Choi, K.Y., Hwang, J.Y., Lee, D.H.: Efficient ID-based group key agreement with bilinear maps. In: Bao, F., Deng, R., Zhou, J. (eds.) PKC 2004. LNCS, vol. 2947, pp. 130–144. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24632-9_10

18. Crockett, E., Peikert, C.: Challenges for ring-LWE. Cryptology ePrint Archive, Report 2016/782 (2016). http://eprint.iacr.org/2016/782

19. Ding, J., Xie, X., Lin, X.: A simple provably secure key exchange scheme based on the learning with errors problem. Cryptology ePrint Archive, Report 2012/688 (2012). http://eprint.iacr.org/2012/688

20. Hoeffding, W.: Probability inequalities for sums of bounded random variables. J. Am. Stat. Assoc. **58**(301), 13–30 (1963)

21. Ingemarsson, I., Tang, D., Wong, C.: A conference key distribution system. IEEE Trans. Inf. Theor. **28**(5), 714–720 (1982)

22. Katz, J., Shin, J.S.: Modeling insider attacks on group key-exchange protocols. In: Proceedings of the 12th ACM Conference on Computer and Communications Security, CCS 2005, pp. 180–189. ACM, New York (2005)

23. Katz, J., Yung, M.: Scalable protocols for authenticated group key exchange. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 110–125. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45146-4_7

24. Katz, J., Yung, M.: Scalable protocols for authenticated group key exchange. J. Cryptol. **20**(1), 85–113 (2007)
25. Kim, Y., Perrig, A., Tsudik, G.: Simple and fault-tolerant key agreement for dynamic collaborative groups. In: Proceedings of the 7th ACM Conference on Computer and Communications Security, CCS 2000, pp. 235–244. ACM, New York (2000)
26. Langlois, A., Stehlé, D., Steinfeld, R.: GGHLite: more efficient multilinear maps from ideal lattices. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 239–256. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_14
27. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_1
28. Peikert, C.: Lattice cryptography for the internet. Cryptology ePrint Archive, Report 2014/070 (2014). http://eprint.iacr.org/2014/070
29. Steer, D.G., Strawczynski, L.: A secure audio teleconference system. In: 21st Century Military Communications - What's Possible?'. Conference Record. Military Communications Conference, MILCOM 1988, October 1988
30. Steiner, M., Tsudik, G., Waidner, M.: Key agreement in dynamic peer groups. IEEE Trans. Parallel Distrib. Syst. **11**(8), 769–780 (2000)
31. Wu, Q., Mu, Y., Susilo, W., Qin, B., Domingo-Ferrer, J.: Asymmetric group key agreement. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 153–170. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01001-9_9
32. Zhang, J., Zhang, Z., Ding, J., Snook, M., Dagdelen, Ö.: Authenticated key exchange from ideal lattices. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 719–751. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_24