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, 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. There has been relatively little work, however, 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 [20, 28, 15, 29, 5, 6, 24, 14, 12, 13, 11, 17, 21, 16, 8, 2, 1, 23, 9, 31] in which N parties wish to agree on a common session key that they can each then use for encrypted 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 [26], 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 [23] 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, 23], 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 $X_i = (z_{i+1}/z_{i-1})^{r_i}$ (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 public 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] Exceptions include the work of Ding et al. [18], which lacks a proof of security; the work of Boneh et al. [10] shows a framework for group key-exchange protocols 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. [23]) 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

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

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 the real session key k is indistinguishable from uniform) to an "unpredictability-based" one (namely, requiring that it is infeasible for an attacker to compute the real session key k). 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 S is a set, then $x_0, x_1, \ldots, x_{\ell-1} \leftarrow S$ denotes uniformly sampling each x_i from S; if χ is a probability distribution, then $x_0, x_1, \ldots, x_{\ell-1} \leftarrow \chi$ denotes independently sampling each x_i according to that distribution. Let $\chi(E)$ denote the probability that event E occurs under distribution χ. We let $\mathrm{Supp}(\chi) = \{x : \chi(x) \neq 0\}$. Given an event E, we let $\overline{E}$ denote its complement. Given a polynomial p_i, let $(p_i)_j$ denote the jth coefficient of p_i. We use $\log(X)$ to denote $\log_2(X)$, and $\exp(X)$ to denote e^X.

We let λ denote a computational security parameter, and ρ a statistical security parameter.

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 uniformly from the ring. More precisely, the Ring-LWE problem is parameterized by (R, q, χ, ℓ) where:

1. $R = \mathbb{Z}[X]/(f(X))$ is a ring, where $f(X)$ is an irreducible polynomial $f(X)$ in the indeterminate X. In this paper, we restrict to the case of $f(X) = X^n + 1$, where n is a power of 2.
2. q is a modulus defining the quotient ring $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$. We restrict to the case where q is prime with $q = 1 \bmod 2n$.
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."
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. In the first case, the samples are generated by choosing $s \leftarrow \chi_s$ and 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 is uniform in R_q and each $e_i \leftarrow \chi_e$ is drawn from the error distribution χ_e. In the second case, each sample (a_i, b_i) is uniformly and independently drawn from $R_q \times R_q$. We let $\mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_s,\chi_e,\ell}(\mathcal{B})$ denote the advantage of algorithm $\mathcal{B}$ in distinguishing these two cases, and define $\mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_s,\chi_e,\ell}(t)$ to be the maximum advantage of any algorithm running in time t. If $\chi = \chi_s = \chi_e$, we write $\mathsf{Adv}_{n,q,\chi,\ell}$ for simplicity.

The noise distribution. The noise distribution $\chi = \chi_s = \chi_e$ is usually a discrete Gaussian distribution on R_q. For power-of-2 cyclotomic rings of the form we consider here, it is possible to sample a polynomial from this distribution by drawing each coefficient of the polynomial independently from the 1-dimensional discrete Gaussian distribution over $\mathbb{Z}_q$ with parameter σ. This distribution, supported on $\{x \in \mathbb{Z}; -q/2 < x < q/2\}$, 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

For two discrete probability distributions P and Q with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$, their *Rényi divergence* is defined as

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

We use the following results (see [30, 26, 25] for proofs):

Proposition 1. *For discrete distributions P and Q with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$ and any f, we have*

$$\mathrm{RD}_2(f(P)||f(Q)) \leq \mathrm{RD}_2(P||Q).$$

Proposition 2. *For discrete distributions P and Q with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$, let $E \subseteq \mathrm{Supp}(Q)$ be an arbitrary event. We have*

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

The second property implies, roughly, that as long as $\mathrm{RD}_2(P\|Q)$ is bounded by some polynomial, then any event E that occurs with negligible probability $Q(E)$ under distribution Q also occurs with negligible probability $P(E)$ under distribution P.

The following theorem bounds the Rényi divergence between the 1-dimensional discrete Gaussian distribution centered at the origin and one centered at a point near the origin.

Theorem 2.1 ([7]). *Fix $m, q, \lambda \in \mathbb{Z}$, a bound $\beta_{\mathsf{Rényi}}$, and σ with $\beta_{\mathsf{Rényi}} < \sigma < q$. Let $e \in \mathbb{Z}$ be such that $|e| \leq \beta_{\mathsf{Rényi}}$. Then*

$$\mathrm{RD}_2((e + D_{\mathbb{Z}_q,\sigma})^m || D^m_{\mathbb{Z}_q,\sigma}) \leq \exp(2\pi m(\beta_{\mathsf{Rényi}}/\sigma)^2).$$

(Here, χ^m denotes m independent samples from distribution χ.)

The above theorem implies that if $\sigma = \Omega(\beta_{\mathsf{Rényi}}\sqrt{m/\log\lambda})$ for some security parameter λ, then $\mathrm{RD}_2((e + D_{\mathbb{Z}_q,\sigma})^m || D^m_{\mathbb{Z}_q,\sigma}) = \mathrm{poly}(\lambda)$.

2.4 Generic Key Reconciliation

In this subsection, we define a generic, one round, two-party key reconciliation mechanism (tailored to the Ring-LWE setting) that allows two parties to derive a shared key if they begin holding "close" ring elements. Formally, a key-reconciliation mechanism KeyRec consists of two algorithms recMsg and recKey, parameterized by a bound β_{Rec} (that may depend on the security parameter). The first algorithm takes as input the security parameter 1^λ and a value $b \in R_q$, and outputs a reconciliation message rec and a key $k \in \{0,1\}^\lambda$. The second algorithm takes as input 1^λ, a value $b' \in R_q$, and rec, and outputs $k' \in \{0,1\}^\lambda$.

Correctness requires that whenever b, b' are "close," then $k' = k$. Specifically, for any b, b' for which each coefficient of $b - b'$ is bounded by β_{Rec}, if we run $(\mathsf{rec}, k) \leftarrow \mathsf{recMsg}(1^\lambda, b)$ followed by $k' := \mathsf{recKey}(1^\lambda, b', \mathsf{rec})$ then $k = k'$.

Security requires that if b is uniform and we derive $(\mathsf{rec}, k) \leftarrow \mathsf{recMsg}(1^\lambda, b)$, then k is computationally indistingiushable from uniform even for an attacker given rec. Formally, the following two distribution ensembles must be computationally indistinguishable:

$$\left\{b \leftarrow R_q; (\mathsf{rec}, k) \leftarrow \mathsf{recMsg}(1^\lambda, b) : (\mathsf{rec}, k)\right\}_{\lambda\in\mathbb{N}},$$

$$\left\{b \leftarrow R_q; (\mathsf{rec}, k) \leftarrow \mathsf{recMsg}(1^\lambda, b); k' \leftarrow \{0,1\}^\lambda : (\mathsf{rec}, k')\right\}_{\lambda\in\mathbb{N}},$$

For some fixed value of λ we denote by $\mathsf{Adv}_{\mathsf{KeyRec}}(\mathcal{B})$ the advantage of adversary $\mathcal{B}$ in distinguishing these distributions, and let $\mathsf{Adv}_{\mathsf{KeyRec}}(t)$ be the maximum advantage of any such adversary running in time t.

Key-reconciliation mechanisms from the literature. The notion of key reconciliation was first introduced by Ding et al. [18], and was later used in several works on two-party key exchange [27, 32, 4]. In the key reconciliation mechanisms of Peikert [27], Zhang et al. [32] and Alkim et al. [4], the agreed-upon key $k = k'$ is close to each of the original values b, b' held by the parties. When instantiating our group key exchange (GKE) protocol with this type of key-reconciliation mechanism, our final GKE protocol is contributory. In other cases [3], the agreed-upon key is determined by the randomness used when running recMsg; instantiating our GKE protocol with this type of key-reconciliation mechanism yields a non-contributory protocol.

3 Group Key Exchange

A group key-exchange protocol allows a session key to be established among $N > 2$ parties. Following prior work [22, 14, 12, 13], we will use the term *group key exchange* (GKE) to denote a protocol secure against a *passive* (eavesdropping) adversary, and 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 [22] presents a compiler that takes any GKE protocol and transforms it into a GAKE protocol. The underlying tool required for this transform is any secure signature scheme; if post-quantum security is needed, then any post-quantum signature scheme can be used. We thus focus our attention on achieving GKE in the remainder of this work.

In the security definition for group key exchange, the adversary observes a single transcript generated by an execution of the protocol. The adversary's goal is then to distinguish the real session key generated in that execution of the protocol from a key that is generated uniformly and independently of that transcript. Formally, given a GKE protocol Π we let $\mathsf{Execute}_{\Pi}(\lambda)$ denote an execution of the protocol (on security parameter λ), resulting in a transcript trans of all messages sent during the course of that execution, along with the session key sk computed by the parties. Protocol Π is secure if the following distribution ensembles are computationally indistinguishable:

$$\{(\mathsf{trans}, \mathsf{sk}) \leftarrow \mathsf{Execute}_{\Pi}(\lambda) : (\mathsf{trans}, \mathsf{sk})\}_{\lambda\in\mathbb{N}},$$
$$\{(\mathsf{trans}, \mathsf{sk}) \leftarrow \mathsf{Execute}_{\Pi}(\lambda), \mathsf{sk}' \leftarrow \{0,1\}^{\lambda} : (\mathsf{trans}, \mathsf{sk}')\}_{\lambda\in\mathbb{N}}.$$

Our protocol Π will be analyzed in the random-oracle model. In this case, fixing some λ, we let $\mathsf{Adv}_{\Pi}^{\mathsf{GKE}}(\mathcal{A})$ denote the advantage of an adversary $\mathcal{A}$ in distinguishing between the distributions above, and define $\mathsf{Adv}_{\Pi}^{\mathsf{GKE}}(t, \mathsf{q})$ to be the maximum advantage of any adversary running in time t and making at most q queries to the random oracle.

4 A Group Key-Exchange Protocol

In this section, we present a group key exchange protocol Π for N parties $P_0, \ldots, P_{N-1}$. Our protocol relies on a key-reconciliation mechanism KeyRec (parameterized by a bound β_{Rec}) as a subroutine.

The overall structure of the protocol is as follows. The first two rounds allow the parties to agree on "close" keys $b_0 \approx \cdots \approx b_{N-1}$. Player $N-1$ then initiates the key-reconciliation mechanism to allow all parties to agree on the same key $k = k_0 = \cdots = k_{N-1} \in \{0,1\}^{\lambda}$. 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 then have each party hash k (using a hash function $\mathcal{H}$) to obtain the final shared key sk.

Our protocol is parameterized by noise distributions $\chi_{\sigma_1}, \chi_{\sigma_2}$, and assumes public parameters $R_q = \mathbb{Z}_q[x]/(x^n + 1)$ along with a uniform value $a \in R_q$. The protocol proceeds as follows:

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 $e_i' \leftarrow \chi_{\sigma_1}$. Each P_i broadcasts $X_i = (z_{i+1} - z_{i-1})s_i + e_i'$.

Round 3: Player P_{N-1} samples $e_{N-1}'' \leftarrow \chi_{\sigma_1}$ and computes

$$b_{N-1} = z_{N-2}Ns_{N-1} + (N-1)\cdot X_{N-1} + (N-2)\cdot X_0 + \cdots + X_{N-3} + e_{N-1}''.$$

It then computes $(\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1})$ and broadcasts rec. Finally, it outputs the session key $\mathsf{sk}_{N-1} = \mathcal{H}(k_{N-1})$.

Key computation: Each player P_i (except P_{N-1}) computes

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

It then sets $k_i = \mathsf{recKey}(b_i, \mathsf{rec})$, and outputs the session key $\mathsf{sk}_i = \mathcal{H}(k_i)$.

The following shows a condition under which each party derives the same session key with all but negligible probability.

Theorem 4.1. *Fix ρ, and assume*

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

Then all parties output the same key except with probability at most $2^{-\rho+1}$.

We refer to Appendix A for the proof.

5 Proof of Security

Here we prove security of our protocol Π. We remark that our proof considers only a classical attacker; in particular, we only allow the attacker 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. *Assume* $2N\sqrt{n}\,\lambda^{3/2}\,\sigma_1^2 + (N-1)\cdot\sigma_1 \leq \beta_{\mathit{R\acute{e}nyi}}$ *and* $\beta_{\mathit{R\acute{e}nyi}} < \sigma_2 < q$, *and model* $\mathcal{H}$ *as a random oracle. Then*

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

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

Proof. Let Expt_0 refer to the experiment in which protocol Π is executed to obtain output $(\mathsf{T}, \mathsf{sk})$, where $\mathsf{T} = (\{z_i\}, \{X_i\}, \mathsf{rec})$ is the transcript of the execution and sk is the final shared session key (more formally, the session key output by P_{N-1}). We also then provide the attacker $\mathcal{A}$ with $(\mathsf{T}, \mathsf{sk})$, and then allow $\mathcal{A}$ to interact with the random oracle used when executing Π. Our goal is to bound the advantage of an attacker in distinguishing between samples $(\mathsf{T}, \mathsf{sk})$ distributed according to Expt_0 and samples $(\mathsf{T}, \mathsf{sk}')$ in which T is distributed the same way but sk' is a uniform key (chosen independently of T). To do so, we show that the probability that $\mathcal{A}$ queries k_{N-1} to the random oracle (which we denote by the event Query) is small; since that is the only way an attacker can distinguish $\mathsf{sk} = \mathcal{H}(k_{N-1})$ from an independent, uniform value, that allows us to prove our desired result. In proving our result, we consider a sequence of experiments, and let $\Pr_i[\cdot]$ denote the probability of an event in Experiment i.

For completeness, we write out the distribution of $(\mathsf{T}, \mathsf{sk})$ in Expt_0:

$$\mathsf{Expt}_0 := \left\{ \begin{array}{ll} a \leftarrow R_q;\ \forall i : s_i, e_i \leftarrow \chi_{\sigma_1}; z_i = as_i + e_i; & \\ e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; e'_0 \leftarrow \chi_{\sigma_2}; & \\ \forall i : X_i = (z_{i+1} - z_{i-1})s_i + e'_i; & \\ e''_{N-1} \leftarrow \chi_{\sigma_1}; & : (\mathsf{T}, \mathsf{sk}) \\ b_{N-1} = e''_{N-1} + z_{N-2}Ns_{N-1} + X_{N-1} \cdot (N-1) + & \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; & \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}) & \end{array} \right\}.$$

Since $\mathsf{Adv}_{\Pi}^{\mathsf{GKE}}(t, \mathsf{q}) \leq \Pr_0[\mathsf{Query}]$, we focus on bounding $\Pr_0[\mathsf{Query}]$ for the rest of the proof.

Experiment 1. In this experiment, X_0 is replaced by $X'_0 = -\sum_{i=1}^{N-1} X_i + e'_0$. The corresponding distribution of $(\mathsf{T}, \mathsf{sk})$ is thus as follows:

$$\mathsf{Expt}_1 := \left\{ \begin{array}{ll} a \leftarrow R_q;\ \forall i : s_i, e_i \leftarrow \chi_{\sigma_1}; z_i = as_i + e_i; & \\ e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; e'_0 \leftarrow \chi_{\sigma_2} & \\ X'_0 = -\displaystyle\sum_{i=1}^{N-1} X_i + e'_0; & \\ \forall i > 0 : X_i = (z_{i+1} - z_{i-1})s_i + e'_i & : (\mathsf{T}, \mathsf{sk}) \\ e''_{N-1} \leftarrow \chi_{\sigma_1}; & \\ b_{N-1} = e''_{N-1} + z_{N-2}Ns_{N-1} + X_{N-1} \cdot (N-1) + & \\ \qquad X'_0 \cdot (N-2) + \cdots + X_{N-3}; & \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X'_0, \ldots, X_{N-1}, \mathsf{rec}) & \end{array} \right\}.$$

The following claim, which is the crux of our proof, relates the probabilities of Query in Expt_0 and Expt_1.

Claim. If $2N\sqrt{n}\,\lambda^{3/2}\,\sigma_1^2 + (N-1) \cdot \sigma_1 \leq \beta_{\mathsf{Rényi}}$, then

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

Proof. Note that we may define the random variables X_0, X'_0 in both experiments Expt_0 and Expt_1. Define the random variable Error (in either experiment) as

$$\mathsf{Error} = \sum_{i=0}^{N-1} (s_i e_{i+1} + s_i e_{i-1}) + \sum_{i=1}^{N-1} e'_i.$$

Defining

$$\mathsf{main} = as_1 s_0 - as_{N-1}s_0 - \mathsf{Error},$$

it is straightforward to verify that

$$\begin{aligned} X_0 &= \mathsf{main} + \mathsf{Error} + e'_0 \\ X'_0 &= \mathsf{main} + e'_0, \end{aligned}$$

where e'_0 is sampled from χ_{σ_2}. Our aim is to apply Theorem 2.1 to show that the Rényi divergence between X_0 and X'_0 (and hence between Expt_0 and Expt_1) is small. To do so, we must first show that the absolute value of each coefficient of Error is bounded by $\beta_{\mathsf{Rényi}}$ with all but negligible probability.

Let $\mathsf{bound}_{\mathsf{Err}}$ be the event that for all j we have $|\mathsf{Error}_j| \leq \beta_{\mathsf{Rényi}}$. Note that

$$|\mathsf{Error}_j| = \left| \left(\sum_{i=0}^{N-1} (s_i e_{i+1} + s_i e_{i-1}) + \sum_{i=1}^{N-1} e'_i \right)_j \right|.$$

Fix $c = \sqrt{\frac{2\lambda}{\pi \log e}}$, and let bound be the event that for all i, j we have $|(e'_0)_j| \leq c\sigma_2$ and $|(s_i)_j|$, $|(e_i)_j|$, $|(e''_{N-1})_j| \leq c\sigma_1$, and that for all $i > 0$ and all j it holds that $|(e'_i)_j| \leq c\sigma_1$. Applying Lemmas A.1 and A.2 (with $\rho = \lambda$), we see that

$$\Pr[\mathsf{bound}] \geq 1 - 2^{-\lambda}$$

and

$$\Pr\left[|(s_i e_j)_v| \geq \sqrt{n}\lambda^{3/2}\sigma_1^2 \mid \mathsf{bound} \right] \leq 2^{-2\lambda+1}.$$

Via a union bound, we thus have

$$\Pr\left[\forall j : \, |\mathsf{Error}_j| \leq 2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1 \mid \mathsf{bound} \right] \geq 1 - 4N \cdot n \cdot 2^{-2\lambda}.$$

Under the assumption that $4Nn \leq 2^{\lambda}$ (which holds for all reasonable settings of the parameters) and using a similar argument as in the proof of Lemma A.2, we conclude that

$$\Pr[\mathsf{bound}_{\mathsf{Err}}] \geq 1 - 2^{-\lambda+1}. \tag{2}$$

When $\mathsf{bound}_{\mathsf{Err}}$ occurs, Theorem 2.1 tells us that

$$\mathrm{RD}_2(\mathsf{Error} + \chi_{\sigma_2} || \chi_{\sigma_2}) \leq \exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2). \tag{3}$$

Therefore,

$$\begin{aligned} \Pr_0[\mathsf{Query}] &\leq \Pr_0[\mathsf{Query} \mid \mathsf{bound}_{\mathsf{Err}}] + \Pr_0[\overline{\mathsf{bound}_{\mathsf{Err}}}] \\ &\leq \Pr_0[\mathsf{Query} \mid \mathsf{bound}_{\mathsf{Err}}] + 2^{-\lambda+1} \\ &\leq \sqrt{\Pr_1[\mathsf{Query} \mid \mathsf{bound}_{\mathsf{Err}}] \cdot \exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2)} + 2^{-\lambda+1} \\ &\leq \sqrt{\Pr_1[\mathsf{Query}] \cdot \frac{\exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2)}{\Pr_1[\mathsf{bound}_{\mathsf{Err}}]}} + 2^{-\lambda+1} \\ &\leq \sqrt{\Pr_1[\mathsf{Query}] \cdot \frac{\exp(2\pi n(\beta_{\mathsf{Rényi}}/\sigma_2)^2)}{1 - 2^{-\lambda+1}}} + 2^{-\lambda+1}. \end{aligned}$$

This completes the proof of the claim. □

In Appendix B, we prove (using arguments similar to those in [23]) that

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

This completes the proof of Theorem 5.1. □

Parameter constraints. Beyond the parameter settings required for hardness of the Ring-LWE problem, the parameters $N, n, \sigma_1, \sigma_2, \lambda, \rho$ of the protocol are also required to satisfy the following:

$$(N^2 + 2N) \cdot \sqrt{n}\rho^{3/2}\sigma_1^2 + (\frac{N^2}{2} + 1)\sigma_1 + (N-2)\sigma_2 \leq \beta_{\mathsf{Rec}} \quad \text{(correctness)} \tag{4}$$

$$2N\sqrt{n}\lambda^{3/2}\sigma_1^2 + (N-1)\sigma_1 \leq \beta_{\mathsf{Rényi}} \quad \text{(security)} \tag{5}$$

$$\sigma_2 = \Omega(\beta_{\mathsf{Rényi}}\sqrt{n/\log\lambda}). \quad \text{(security)} \tag{6}$$

Thus, fixing the ring, the noise distributions, and the security parameters λ, ρ induces a bound on the maximum number of parties the protocol can support.

6 Acknowledgments

This material is based on work performed under financial assistance award 70NANB15H328 from the U.S. Department of Commerce, National Institute of Standards and Technology. Work by Dana Dachman-Soled was additionally supported in part by NSF grants #CNS-1840893 and #CNS-1453045, and by a research partnership award from Cisco.

References

1. Michel Abdalla, Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. Password-based group key exchange in a constant number of rounds. In *9th Intl. Conference on Theory and Practice of Public Key Cryptography (PKC)*, volume 3958 of *Lecture Notes in Computer Science*, pages 427–442. Springer, 2006.
2. Michel Abdalla and David Pointcheval. A scalable password-based group key exchange protocol in the standard model. In *Advances in Cryptology—Asiacrypt 2006*, volume 4284 of *Lecture Notes in Computer Science*, pages 332–347. Springer, 2006.
3. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. NewHope without reconciliation. Cryptology ePrint Archive, Report 2016/1157, 2016. http://eprint.iacr.org/2016/1157.
4. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange—a new hope. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 327–343, Austin, TX, 2016. USENIX Association.
5. Klaus Becker and Uta Wille. Communication complexity of group key distribution. In *Proceedings of the 5th ACM Conference on Computer and Communications Security*, CCS '98, pages 1–6, New York, NY, USA, 1998.
6. Mihir Bellare and Phillip Rogaway. Provably secure session key distribution: The three party case. In *27th Annual ACM Symposium on Theory of Computing*, pages 57–66, Las Vegas, NV, USA, May 29 – June 1, 1995. ACM Press.

7. Andrej Bogdanov, Siyao Guo, Daniel Masny, Silas Richelson, and Alon Rosen. On the hardness of learning with rounding over small modulus. In *Theory of Cryptography Conference*, pages 209–224. Springer, 2016.
8. Jens-Matthias Bohli, Maria Isabel Gonzalez Vasco, and Rainer Steinwandt. 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. Jens-Matthias Bohli, María Isabel González Vasco, and Rainer Steinwandt. Secure group key establishment revisited. *International Journal of Information Security*, 6(4):243–254, Jul 2007.
10. Dan Boneh, Darren Glass, Daniel Krashen, Kristin Lauter, Shahed Sharif, Alice Silverberg, Mehdi Tibouchi, and Mark Zhandry. Multiparty non-interactive key exchange and more from isogenies on elliptic curves. *arXiv preprint arXiv:1807.03038*, 2018.
11. Emmanuel Bresson and Dario Catalano. Constant round authenticated group key agreement via distributed computation. In Feng Bao, Robert Deng, and Jianying Zhou, editors, *PKC 2004: 7th Intl. Workshop on Theory and Practice in Public Key Cryptography*, volume 2947 of *Lecture Notes in Computer Science*, pages 115–129, Singapore, March 1–4, 2004. Springer.
12. Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. Provably authenticated group Diffie-Hellman key exchange – the dynamic case. In Colin Boyd, editor, *Advances in Cryptology—Asiacrypt 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 290–309, Gold Coast, Australia, December 9–13, 2001. Springer.
13. Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. Dynamic group Diffie-Hellman key exchange under standard assumptions. In Lars R. Knudsen, editor, *Advances in Cryptology—Eurocrypt 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 321–336, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer.
14. Emmanuel Bresson, Olivier Chevassut, David Pointcheval, and Jean-Jacques Quisquater. Provably authenticated group Diffie-Hellman key exchange. In *ACM CCS 01: 8th Conference on Computer and Communications Security*, pages 255–264, Philadelphia, PA, USA, November 5–8, 2001. ACM Press.
15. Mike Burmester and Yvo Desmedt. A secure and efficient conference key distribution system (extended abstract). In Alfredo De Santis, editor, *Advances in Cryptology—Eurocrypt'94*, volume 950 of *Lecture Notes in Computer Science*, pages 275–286. Springer, 1995.
16. Mike Burmester and Yvo Desmedt. A secure and scalable group key exchange system. *Information Processing Letters*, 94(3):137–143, May 2005.
17. Kyu Young Choi, Jung Yeon Hwang, and Dong Hoon Lee. Efficient ID-based group key agreement with bilinear maps. In Feng Bao, Robert Deng, and Jianying Zhou, editors, *PKC 2004: 7th Intl. Workshop on Theory and Practice in Public Key Cryptography*, volume 2947 of *Lecture Notes in Computer Science*, pages 130–144, Singapore, March 1–4, 2004. Springer.
18. Jintai Ding, Xiang Xie, and Xiaodong Lin. 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.
19. Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
20. I. Ingemarsson, D. Tang, and C. Wong. A conference key distribution system. *IEEE Trans. Inf. Theor.*, 28(5):714–720, September 1982.

21. Jonathan Katz and Ji Sun Shin. Modeling insider attacks on group key-exchange protocols. In *Proceedings of the 12th ACM Conference on Computer and Communications Security*, CCS '05, pages 180–189, New York, NY, USA, 2005. ACM.
22. Jonathan Katz and Moti Yung. Scalable protocols for authenticated group key exchange. In Dan Boneh, editor, *Advances in Cryptology—Crypto 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 110–125, Santa Barbara, CA, USA, 2003. Springer.
23. Jonathan Katz and Moti Yung. Scalable protocols for authenticated group key exchange. *Journal of Cryptology*, 20(1):85–113, 2007.
24. Yongdae Kim, Adrian Perrig, and Gene Tsudik. Simple and fault-tolerant key agreement for dynamic collaborative groups. In *Proceedings of the 7th ACM Conference on Computer and Communications Security*, CCS '00, pages 235–244, New York, NY, USA, 2000.
25. Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology—Eurocrypt 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 239–256, Copenhagen, Denmark, May 11–15, 2014. Springer.
26. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology—Eurocrypt 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23, French Riviera, May 30 – June 3, 2010. Springer.
27. Chris Peikert. Lattice cryptography for the internet. Cryptology ePrint Archive, Report 2014/070, 2014. http://eprint.iacr.org/2014/070.
28. D. G. Steer and L. Strawczynski. A secure audio teleconference system. In *MILCOM 88, 21st Century Military Communications - What's Possible?'. Conference record. Military Communications Conference*, Oct 1988.
29. M. Steiner, G. Tsudik, and M. Waidner. Key agreement in dynamic peer groups. *IEEE Transactions on Parallel and Distributed Systems*, 11(8):769–780, Aug 2000.
30. Tim Van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
31. Qianhong Wu, Yi Mu, Willy Susilo, Bo Qin, and Josep Domingo-Ferrer. Asymmetric group key agreement. In Antoine Joux, editor, *Advances in Cryptology—Eurocrypt 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 153–170, Cologne, Germany, April 26–30, 2009. Springer.
32. Jiang Zhang, Zhenfeng Zhang, Jintai Ding, Michael Snook, and Özgür Dagdelen. Authenticated key exchange from ideal lattices. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology—Eurocrypt 2015, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 719–751, Sofia, Bulgaria, April 26–30, 2015. Springer.

A Proof of Correctness

Theorem 4.1 (Restated)**.** *Fix ρ, and assume*

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

Then all parties output the same key except with probability at most $2^{-\rho+1}$.

Proof. We begin by introducing the following lemmas to analyze probabilities that each coordinate of $s_i, e_i, e'_i, e''_{N-1}, e'_0$ are "short" for all i, and conditioned on the first event, $s_i e_i$ is "short".

Lemma A.1. *Given $s_i, e_i, e'_i, e''_{N-1}, e'_0$ for all i as defined in the group key exchange protocol, fix $c = \sqrt{\frac{2\rho}{\pi \log e}}$, and let bound_ρ denote the event that for all i and all coordinate indices j, $|(e'_0)_j| \leq c\sigma_2$ and $|(s_i)_j|$, $|(e_i)_j|$, $|(e''_{N-1})_j| \leq c\sigma_1$, and that for all $i > 0$ and all j it holds that $|(e'_i)_j| \leq c\sigma_1$, we have*

$$\Pr[\mathsf{bound}_\rho] \geq 1 - 2^{-\rho}.$$

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

$$\begin{aligned}\Pr[|v| \geq c\sigma + 1; v \leftarrow D_{\mathbb{Z}_q,\sigma}] \leq 2 \sum_{x=\lfloor c\sigma+1 \rceil}^{\infty} D_{\mathbb{Z}_q,\sigma}(x) &\leq \frac{2}{\sigma} \int_{c\sigma}^{\infty} e^{-\frac{\pi x^2}{\sigma^2}} dx \\ &= \frac{2}{\sqrt{\pi}} \int_{\frac{\sqrt{\pi}}{\sigma}(c\sigma)}^{\infty} e^{-t^2} dt \leq e^{-c^2\pi}.\end{aligned}$$

Note that there are $3nN$ coordinates sampled from distribution $D_{\mathbb{Z}_q,\sigma_1}$, and n coordinates sampled from distribution $D_{\mathbb{Z}_q,\sigma_2}$ in total. Under the assumption that $3nN + n \leq e^{c^2\pi/2}$ (which holds for all reasonable settings for the parameters), we have:

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

□

Lemma A.2. *Given bound_ρ as defined in Lemma A.1, let $\mathsf{product}_{\mathsf{s_i},\mathsf{e_j}}$ denote the event that, for all v, $|(s_i e_j)_v| \leq \sqrt{n}\rho^{3/2}\sigma_1^2$,*

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

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 $c_v = \sum_{u=0}^{n-1} (s_i)_u (e_j)^*_{v-u}$. If $v - u \geq 0$, $(e_j)^*_{v-u} = (e_j)_{v-u}$. $(e_j)^*_{v-u} = -(e_j)_{v-u+n}$ 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 [19], we derive

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

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 fixing $\delta = \sqrt{n}\rho^{3/2}\sigma_1^2$, we obtain

$$\Pr[|(s_ie_j)_v| \geq \sqrt{n}\rho^{3/2}\sigma_1^2 \mid \mathsf{bound}_\rho] \leq 2^{-2\rho+1}. \tag{7}$$

Finally, via a union bound, we thus have

$$\Pr[\mathsf{product}_{\mathsf{s_i},\mathsf{e_j}}|\mathsf{bound}_\rho] = \Pr[\forall v : |(s_ie_j)_v| \leq \sqrt{n}\rho^{3/2}\sigma_1^2] \geq 1 - 2n \cdot 2^{-2\rho}. \tag{8}$$

□

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. Formally, if for all i and j that the j^{th} coefficient of $|b_{N-1} - b_i| \leq \beta_{\mathsf{Rec}}$, then for all i, $k_i = k_{N-1}$.

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

$$= \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & \ldots & 0 & -1 \\ -1 & 1 & 0 & 0 & \ldots & 0 & 0 \\ 0 & -1 & 1 & 0 & \ldots & 0 & 0 \\ 0 & 0 & -1 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & & \\ 0 & 0 & 0 & 0 & \ldots & -1 & 1 \end{bmatrix}}_{\boldsymbol{M}} \underbrace{\begin{bmatrix} as_0s_1 \\ as_1s_2 \\ as_2s_3 \\ as_3s_4 \\ \vdots \\ as_{N-2}s_{N-1} \\ as_{N-1}s_0 \end{bmatrix}}_{\boldsymbol{S}} + \underbrace{\begin{bmatrix} s_0e_1 - s_0e_{N-1} + e'_0 \\ s_1e_2 - s_1e_0 + e'_1 \\ s_2e_3 - s_2e_1 + e'_2 \\ s_3e_4 - s_3e_2 + e'_3 \\ \vdots \\ s_{N-2}e_{N-3} - s_{N-2}e_{N-3} + e'_{N-2} \\ s_{N-1}e_0 - s_{N-1}e_{N-2} + e'_{N-1} \end{bmatrix}}_{\boldsymbol{E}}. \tag{9}$$

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

$$\begin{aligned} b_{N-1} - b_i &= (\boldsymbol{B}_{N-1} - \boldsymbol{B}_i)\,\boldsymbol{E} + N(s_{N-1}e_{N-2} - s_ie_{i-1}) + e''_{N-1} \\ &= (N-i-1) \cdot \left(\sum_{\substack{j \in \mathbb{Z} \cap [0, i-1] \\ \text{and } j = N-1}} s_je_{j+1} - s_je_{j-1} + e'_j \right) + e''_{N-1} \\ &\quad + (-i-1) \left(\sum_{j=i}^{N-2} s_je_{j+1} - s_je_{j-1} + e'_j \right) + N(s_{N-1}e_{N-2} - s_ie_{i-1}) \end{aligned}$$

Observe that for an arbitrary $i \in [N]$, and in any coordinate of the sum above, 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 sampled from χ_{σ_1}, at most $N-2$ terms of e'_0 sampled from χ_{σ_2}, and one term of e''_{N-1}.

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 $\sqrt{n}\rho^{3/2}\sigma_1^2$. Under that assumption that assuming $2n(N^2+2N) \leq 2^{\rho}$ (which holds for all reasonable settings of the parameters) and using a union bound, it is straightforward to see

$$\Pr[\overline{\mathsf{product}_{\mathsf{ALL}}}|\mathsf{bound}_{\rho}] \leq (N^2+2N) \cdot 2n2^{-2\rho} \leq 2^{-\rho}.$$

Let fail 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

$$\begin{aligned}\Pr[\mathsf{fail}] &= \Pr[\mathsf{fail}|\mathsf{bound}_{\rho}] \cdot \Pr[\mathsf{bound}_{\rho}] + \Pr[\mathsf{fail}|\overline{\mathsf{bound}_{\rho}}] \cdot \Pr[\overline{\mathsf{bound}_{\rho}}] && (10)\\ &\leq \Pr[\overline{\mathsf{product}_{\mathsf{ALL}}}] \cdot 1 + 1 \cdot \Pr[\overline{\mathsf{bound}_{\rho}}] \leq 2 \cdot 2^{-\rho}, && (11)\end{aligned}$$

which completes the proof.

□

B Concluding the Proof of Security

Theorem 5.1 (Restated)**.** *Assume $2N\sqrt{n}\,\lambda^{3/2}\,\sigma_1^2 + (N-1)\cdot\sigma_1 \leq \beta_{\mathsf{Rényi}}$ and $\beta_{\mathsf{Rényi}} < \sigma_2 < q$, and model $\mathcal{H}$ as a random oracle. Then*

$$\begin{aligned}&\mathsf{Adv}_{\Pi}^{\mathsf{GKE}}(t,\mathsf{q}) \leq 2^{-\lambda+1}\\ &\quad + \sqrt{\left(N \cdot \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{\mathsf{q}}{2^{\lambda}}\right) \cdot \frac{\exp\left(2\pi n\left(\beta_{\mathsf{Rényi}}/\sigma_2\right)^2\right)}{1-2^{-\lambda+1}}},\end{aligned}$$

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

Proof. (Continued) Recall that Experiment 0 is the real world experiment. We have that $\mathsf{Adv}_{\Pi}^{\mathsf{GKE}}(t,\mathsf{q}) \leq \Pr_0[\mathsf{Query}]$, where Query is the event that k_{N-1} is among the adversary $\mathcal{A}$'s random oracle queries and $\Pr_i[\mathsf{Query}]$ is the probability that event Query happens in Experiment i.

In Experiment 1, we switched from X_0 as sampled in the real world to $X'_0 = -\sum_{i=1}^{N-1} X_i + e'_0$ and showed (see Equation 1) that

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

Therefore, to prove the theorem, it remains to show that

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

We do so by considering a sequence of experiments as follows:

Experiment 2. In this experiment, z_0 is replaced by a uniform element in R_q. The corresponding distribution of $(\mathsf{T}, \mathsf{sk})$ is thus as follows:

$$\mathsf{Expt}_2 := \left\{ \begin{array}{l} a, z_0 \leftarrow R_q;\ \forall i \geq 1 : s_i, e_i \leftarrow \chi_{\sigma_1}; z_i = as_i + e_i; \\ e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; e'_0 \leftarrow \chi_{\sigma_2} \\ X_0 = -\displaystyle\sum_{i=1}^{N-1} X_i + e'_0, \forall i \geq 1 : X_i = (z_{i+1} - z_{i-1})s_i + e'_i \qquad : (\mathsf{T}, \mathsf{sk}) \\ e''_{N-1} \leftarrow \chi_{\sigma_1}; \\ b_{N-1} = e''_{N-1} + z_{N-2}Ns_{N-1} + X_{N-1} \cdot (N-1) + \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}). \end{array} \right\}.$$

Claim. For any algorithm $\mathcal{A}$ running in time t, we have

$$|\Pr_2[\mathsf{Query}] - \Pr_1[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1), \tag{12}$$

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

Proof. We first consider an experiment Expt'_1 which is identical to Expt_1 except for (a, z_0) given as input. For algorithm $\mathcal{A}$ running in time t, let $\mathcal{B}$ be an algorithm running in time t_1 which takes as input (a, z_0), generates $(\mathsf{T}, \mathsf{sk})$ according to Expt'_1, runs $\mathcal{A}(\mathsf{T}, \mathsf{sk})$ as a subroutine and outputs whatever $\mathcal{A}$ outputs. t_1 is then equal to t plus a minor overhead for the simulation of the security experiment for $\mathcal{A}$.

It is straightforward to see that if (a, z_0) is sampled from $A_{n,q,\chi_{\sigma_1}}$, then Expt'_1 is identical to Expt_1, and if (a, z_0) is sampled from R_q^2, Expt'_1 is identical to Expt_2.

Therefore the difference of algorithm $\mathcal{A}$'s success probability in Experiment 1 and Experiment 2 is bounded by probability that $\mathcal{B}$ running in time t_1 distinguishes $A_{n,q,\chi_{\sigma_1}}$ from R_q^2 given one sample. Since

$$\mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) \geq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},2}(t_1) \geq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},1}(t_1),$$

for simplicity, we conclude that:

$$|\Pr_2[\mathsf{Query}] - \Pr_1[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1), \tag{13}$$

□

Recall that in the previous experiment, we switched z_0 to be uniformly distributed in R_q. In next two experiments, we switch z_1, X_1 to be elements uniformly distributed in R_q.

Experiment 3. In this experiment, z_0 is replaced by $z_2 - r_1$, and X_1 is replaced by $r_1 s_1 + e'_1$, where r_1 is uniform in R_q. The corresponding distribution of $(\mathsf{T}, \mathsf{sk})$ is thus as follows:

Since r_1 is uniform, then $z_2 - r_1$ is also uniform. Thus, we conclude that Experiment 3 is identical to Experiment 2 up to variable substitution, namely

$$\Pr_3[\mathsf{Query}] = \Pr_2[\mathsf{Query}]. \tag{14}$$

$$\mathsf{Expt}_3 := \left\{ \begin{array}{ll} a, r_1 \leftarrow R_q;\ \forall i \geq 1 : s_i, e_i \leftarrow \chi_{\sigma_1}; z_i = as_i + e_i; & \\ z_0 = z_2 - r_1; & \\ \forall i \geq 1 : e'_i \leftarrow \chi_{\sigma_1};\ e'_0 \leftarrow \chi_{\sigma_2}; & \\ X_0 = -\sum_{i=1}^{N-1} X_i + e'_0; X_1 = r_1 s_1 + e'_1; & : (\mathsf{T}, \mathsf{sk}) \\ \forall i \geq 2 : X_i = (z_{i+1} - z_{i-1}) s_i + e'_i; & \\ e''_{N-1} \leftarrow \chi_{\sigma_1}; & \\ b_{N-1} = e''_{N-1} + z_{N-2} N s_{N-1} + X_{N-1} \cdot (N-1) + & \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; & \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}). & \end{array} \right\}.$$

Experiment 4. In this experiment, z_1, X_1 are replaced by uniform elements in R_q. The corresponding distribution of $(\mathsf{T}, \mathsf{sk})$ is thus as follows:

$$\mathsf{Expt}_4 := \left\{ \begin{array}{ll} a, r_1 \leftarrow R_q;\ \forall i \geq 2 : s_i, e_i \leftarrow \chi_{\sigma_1}; z_i = as_i + e_i; & \\ z_0 = z_2 - r_1, z_1 \leftarrow R_q; & \\ e'_2, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; e'_0 \leftarrow \chi_{\sigma_2}; & \\ X_0 = -\sum_{i=1}^{N-1} X_i + e'_0, X_1 \leftarrow R_q; & \\ \forall i \geq 2 : X_i = (z_{i+1} - z_{i-1}) s_i + e'_i, & : (\mathsf{T}, \mathsf{sk}) \\ e''_{N-1} \leftarrow \chi_{\sigma_1}; & \\ b_{N-1} = e''_{N-1} + z_{N-2} N s_{N-1} + X_{N-1} \cdot (N-1) + & \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; & \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}). & \end{array} \right\}.$$

Claim. For any algorithm $\mathcal{A}$ running in time t, we have

$$|\Pr_4[\mathsf{Query}] - \Pr_3[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1), \tag{15}$$

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

Proof. We first consider an experiment Expt_3' which is identical to Expt_3 except for (a, z_1), (r_1, X_1) given as input. For algorithm $\mathcal{A}$ running in time t, let $\mathcal{B}$ be an algorithm running in time t_1 that takes as input (a, z_1), (r_1, X_1), generates $(\mathsf{T}, \mathsf{sk})$ according to Expt_3'. $\mathcal{B}$ then runs $\mathcal{A}(\mathsf{T}, \mathsf{sk})$ as a subroutine and outputs whatever $\mathcal{A}$ outputs. t_1 is then equal to t plus a minor overhead for the simulation of the security experiment for $\mathcal{A}$.

It is clear to see that if (a, z_1) and (r_1, X_1) are sampled from $A_{n,q,\chi_{\sigma_1}}$, then Expt_3' is identical to Expt_3. If (a, z_1) and (r_1, X_1) are sampled from $\mathcal{U}(R_q^2)$, Expt_3' is identical to Expt_4.

Therefore the difference of algorithm $\mathcal{A}$ successful probability in Experiment 3 and Experiment 4 is bounded by the advantage of adversary $\mathcal{B}$ running in time t_1 in distinguishing $A_{n,q,\chi_{\sigma_1}}$ from $\mathcal{U}(R_q^2)$ given two samples. Thus, we conclude

$$|\Pr_4[\mathsf{Query}] - \Pr_3[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1). \tag{16}$$

□

Experiment 5. In this experiment, z_0 is replaced by a uniform element in R_q. The corresponding distribution is denoted as Expt_5. We leave the formal definition of Expt_5 implicit for simplicity

It is easy to see that the corresponding distribution Expt_5 is identical to Expt_4 by substituting variable z_0 for $z_2 - r_1$. Thus,

$$\Pr_5[\mathsf{Query}] = \Pr_4[\mathsf{Query}]. \tag{17}$$

In the case that $N \geq 3$, we present the following sequence of experiments from Experiment 6 to Experiment $3N-4$. For $i = 2, 3, \ldots, N-2$, we define three experiments Experiment $3i$, Experiment $3i+1$, Experiment $3i+2$. It is ensured that in the experiments prior to Experiment $3i$, we already switched z_j, X_j for all $0 \leq j \leq i-1$. In Experiment $3i$, Experiment $3i+1$ and Experiment $3i+2$, we replace z_i and X_i by random elements in R_q. Experiment $3i$, Experiment $3i+1$, Experiment $3i+2$ are formally defined as follows:

Experiment $3i$. The experiment proceeds exactly the same as Experiment $3i-1$, except for setting $z_{i-1} = z_{i+1} - r_i, X_i = r_i s_i + e_i'$, where r_1 is uniform in R_q. The corresponding distribution of $(\mathsf{T}, \mathsf{sk})$ is thus as follows, denoted Expt_{3i}:

Experiment $3i+1$. In this experiment, z_i, X_i are replaced by uniform elements in R_q. The corresponding distribution of $(\mathsf{T}, \mathsf{sk})$ is thus as follows, denoted Expt_{3i+1}:

Experiment $3i+2$. In this experiment, z_{i-1} is replaced by a uniform element in R_q. The corresponding distribution is denoted as Expt_{3i+2}. We leave the formal definition of Expt_{3i+2} implicit for simplicity.

$$\mathsf{Expt}_{3i} := \left\{ \begin{array}{ll} a, r_i \leftarrow R_q;\ \forall j \geq i : s_j, e_j \leftarrow \chi_{\sigma_1}; z_j = as_j + e_j; & \\ z_0, \ldots, z_{i-2} \leftarrow R_q, z_{i-1} = z_{i+1} - r_i; & \\ e'_i, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}, e'_0 \leftarrow \chi_{\sigma_2}; & \\ X_0 = -\sum_{i=1}^{N-1} X_i + e'_0, X_1, \ldots, X_{i-1} \leftarrow R_q; & : (\mathsf{T}, \mathsf{sk}) \\ X_i = r_i s_i + e'_i;\ \forall j \geq i : X_{j+1} = (z_{j+2} - z_j)s_{j+1} + e'_{j+1} & \\ e''_{N-1} \leftarrow \chi_{\sigma_1}; & \\ b_{N-1} = e''_{N-1} + z_{N-2}Ns_{N-1} + X_{N-1} \cdot (N-1) + & \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; & \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}). & \end{array} \right\}.$$

$$\mathsf{Expt}_{3i+1} := \left\{ \begin{array}{ll} a, r_i \leftarrow R_q;\ \forall j \geq i+1 : s_j, e_j \leftarrow \chi_{\sigma_1}; z_j = as_j + e_j; & \\ z_0, \ldots, z_{i-2} \leftarrow R_q, z_{i-1} = z_{i+1} - r_i, z_i \leftarrow R_q, & \\ e'_1, \ldots, e'_{N-1} \leftarrow \chi_{\sigma_1}; e'_0 \leftarrow \chi_{\sigma_2} & \\ X_0 = -\sum_{i=1}^{N-1} X_i + e'_0, X_1, \ldots, X_i \leftarrow R_q, & : (\mathsf{T}, \mathsf{sk}) \\ \forall j \geq i+1, X_j = (z_{j+1} - z_{j_1})s_j + e'_j; & \\ e''_{N-1} \leftarrow \chi_{\sigma_1}; & \\ b_{N-1} = e''_{N-1} + z_{N-2}Ns_{N-1} + X_{N-1} \cdot (N-1) + & \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; & \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}). & \end{array} \right\}.$$

Using similar arguments as proving (in)equalities (14), (15) and (17), we conclude that:

$$\Pr_{3i}[\mathsf{Query}] = \Pr_{3i-1}[\mathsf{Query}]; \tag{18}$$

$$|\Pr_{3i+1}[\mathsf{Query}] - \Pr_{3i}[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1); \tag{19}$$

$$\Pr_{3i+2}[\mathsf{Query}] = \Pr_{3i+1}[\mathsf{Query}]; \tag{20}$$

Note that in Experiment $3N-4$, the last experiment of the experiment sequence above, we already switched all the z_i, X_i up to z_{N-1}, X_{N-1}. We construct the next two experiments to switch $z_{N-1}, X_{N-1}, b_{N-1}$.

Experiment $3N-3$. The experiment proceeds exactly the same as Experiment $3N-4$, except for setting $z_{N-2} = r_2, X_{N-1} = r_1 s_{N-1} + e'_{N-1}, z_0 = r_1 + r_2$, where r_1, r_2 are uniform in R_q. The corresponding distribution is thus as follows:

Since r_1, r_2 are uniform, $r_1 + r_2$ is then also uniform. Thus we conclude that Experiment $3N-3$ is identical to Experiment $3N-4$ up to variable substitution, namely,

$$\Pr_{3N-3}[\mathsf{Query}] = \Pr_{3N-4}[\mathsf{Query}]; \tag{21}$$

$$\mathsf{Expt}_{3N-3} := \left\{ \begin{array}{ll} a, r_1, r_2 \leftarrow R_q, s_{N-1}, e_{N-1} \leftarrow \chi_{\sigma_1}; z_0 = r_1 + r_2, & \\ z_1, \ldots, z_{N-3} \leftarrow R_q, z_{N-2} = r_2, & \\ z_{N-1} = as_{N-1} + e_{N-1}; e'_0 \leftarrow \chi_{\sigma_2}; e'_{N-1} \leftarrow \chi_{\sigma_1}; & \\ X_0 = -\sum_{i=1}^{N-1} X_i + e'_0, X_1, \ldots, X_{N-2} \leftarrow R_q, & \\ X_{N-1} = r_1 s_{N-1} + e'_{N-1}; e''_{N-1} \leftarrow \chi_{\sigma_1}; & : (\mathsf{T}, \mathsf{sk}) \\ b_{N-1} = e''_{N-1} + r_2 N s_{N-1} + X_{N-1} \cdot (N-1) + & \\ \qquad X_0 \cdot (N-2) + \cdots + X_{N-3}; & \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}). & \end{array} \right\}.$$

Experiment $3N-2$**.** In this experiment, $z_{N-1}, X_{N-1}, b_{N-1}$ are replaced by uniform elements in R_q. The corresponding distribution is thus as follows: :

$$\mathsf{Expt}_{3N-2} := \left\{ \begin{array}{ll} a \leftarrow R_q; \forall i : z_i \leftarrow R_q; & \\ e'_0 \leftarrow \chi_{\sigma_2}; r_1, r_2 \leftarrow R_q & \\ X_0 = -\sum_{i=1}^{N-1} X_i + e'_0, X_1, \ldots, X_{N-1} \leftarrow R_q & : (\mathsf{T}, \mathsf{sk}) \\ b_{N-1} \leftarrow R_q; & \\ (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}); \mathsf{sk} = \mathcal{H}(k_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}). & \end{array} \right\}.$$

Claim. For any algorithm $\mathcal{A}$ running in time t, we have

$$|\Pr_{3N-2}[\mathsf{Query}] - \Pr_{3N-3}[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1), \tag{22}$$

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

Proof. Since r_2 is uniform in R_q and N is invertible over R_q, then $r_2 N$ is uniformly distributed in R_q. It is easy to see that $(s_{N-1},\ r_2 N s_{N-1} + e''_{N-1})$ forms an RLWE instance. We let $b_{\mathsf{RLWE}} = r_2 N s_{N-1} + e''_{N-1}$.

We consider an experiment Expt'_{3N-3} which is identical to Expt_{3N-3} except for (a, z_{N-1}), (r_1, X_{N-1}), and $(r_2 N, b_{\mathsf{RLWE}})$ given as input. Given an algorithm $\mathcal{A}$ running in time t, let $\mathcal{B}$ be an algorithm that takes as input (a, z_{N-1}), (r_1, X_{N-1}), and $(r_2 N, b_{\mathsf{RLWE}})$, generates $(\mathsf{T}, \mathsf{sk})$ according to Expt'_{3N-3}. $\mathcal{B}$ runs $\mathcal{A}(\mathsf{T}, \mathsf{sk})$ as a subroutine and outputs whatever $\mathcal{A}$ outputs. Running time t_1 of $\mathcal{B}$ then equals to t plus a minor overhead for the simulation of the security experiment for $\mathcal{A}$.

It is straightforward to see that if (a, z_{N-1}), (r_1, X_1), and $(r_2 N, b_{\mathsf{RLWE}})$ are sampled from $A_{n,q,\chi_{\sigma_1}}$, then Expt'_{3N-3} is identical to Expt_{3N-3}. If (a, z_{N-1}), (r_1, X_{N-1}), and $(r_2 N, b_{\mathsf{RLWE}})$ are sampled from R_q^2, then Expt'_{3N-3} is identical to Expt_{3N-2}, since when b_{RLWE} is sampled uniformly at random, $b_{\mathsf{RLWE}} + X_{N-1} \cdot (N-1) + X_0 \cdot (N-2) + \cdots + X_{N-3}$ is also uniformly distributed over R_q.

Therefore the difference of algorithm $\mathcal{A}$'s success probability in Experiment 3N - 2 and Experiment 3N - 3 is bounded by the advantage of adversary $\mathcal{B}$ running in time t_1 in distinguishing Ring-LWE from R_q given three samples. Thus, we conclude that

$$|\Pr_{3N-2}[\mathsf{Query}] - \Pr_{3N-3}[\mathsf{Query}]| \leq \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1), \tag{23}$$

□

Experiment $3N-1$**.** In this experiment, k_{N-1} is replaced by random element in $\{0,1\}^\lambda$. The corresponding distribution is thus as follows:

$$\mathsf{Expt}_{\mathsf{final}} := \left\{ \begin{array}{ll} a \leftarrow R_q; z_0, \ldots, z_{N-1} \leftarrow R_q; e'_0 \leftarrow \chi_{\sigma_1}; & \\ X_0 = -\displaystyle\sum_{i=1}^{N-1} X_i + e'_0, X_1, \ldots, X_{N-1} \leftarrow R_q & \\ b_{N-1} \leftarrow R_q; (\mathsf{rec}, k_{N-1}) = \mathsf{recMsg}(b_{N-1}) & : (\mathsf{T}, \mathsf{sk}) \\ k'_{N-1} \leftarrow \{0,1\}^\lambda; \mathsf{sk} = \mathcal{H}(k'_{N-1}); & \\ \mathsf{T} = (z_0, \ldots, z_{N-1}, X_0, \ldots, X_{N-1}, \mathsf{rec}); & \end{array} \right\}.$$

Given transcript T, and b_{N-1} which is uniformly distributed, using a straight forward reduction, we obtain advantage of adversary $\mathcal{B}$ running in time t_2 in distinguishing k_{N-1} computed by $\mathsf{recMsg}(b_{N-1})$ from a uniform bit string k'_{N-1} with length λ is at least $|\Pr_{3N-1}[\mathsf{Query}] - \Pr_{3N-2}[\mathsf{Query}]|$, namely,

$$|\Pr_{3N-1}[\mathsf{Query}] - \Pr_{3N-2}[\mathsf{Query}]| \leq \mathsf{Adv}_{\mathsf{KeyRec}}(t_2). \tag{24}$$

Note that t_2 equals to the running time of adversary $\mathcal{A}$ attacking the protocol $\mathit{\Pi}$, plus a minor overhead for simulating experiment for $\mathcal{A}$.

Finally, since adversary attacking the GKE protocol $\mathit{\Pi}$ makes at most q queries to the random oracle, $\Pr_{3N-1}[\mathsf{Query}] = \frac{\mathsf{q}}{2^\lambda} \in \mathsf{negl}(\lambda)$. Combining Equations (12) - (24), we have

$$\Pr_1[\mathsf{Query}] \leq N \cdot \mathsf{Adv}^{\mathsf{RLWE}}_{n,q,\chi_{\sigma_1},3}(t_1) + \mathsf{Adv}_{\mathsf{KeyRec}}(t_2) + \frac{\mathsf{q}}{2^\lambda}. \tag{25}$$

The theorem now follows immediately from Equations (1), and (25). □