EQSIGN: Practical Digital Signatures from the Non-Abelian Hidden Subgroup Problem and Information Theoretic Equivocation

Samuel Lavery

Trustless Privacy Inc.
sam@trustlessprivacy.com
Signal:samlavery.07

January 23, 2025

Abstract

We present a novel digital signature scheme grounded in non-commutative cryptography and implemented using a bilinear matrix group platform. At the core of our design is a unique equivocation function that obfuscates intermediate elements, effectively concealing outputs and minimizing observable information leakage. To the best of our knowledge, this is the first digital signature scheme to combine information-theoretic security with computational hardness, relying on a challenging instance of the Non-Abelian Hidden Subgroup Problem (NAHSP) and strengthened by provable information theoretic guarantees. This dual-layered security approach ensures robustness against both classical and quantum adversaries while maintaining communication overheads competitive with RSA. Our work represents a significant advancement toward efficient, quantum-resilient digital signatures for real-world applications. This paper is an early pre-release intended to invite collaboration and feedback. The work is presented for research purposes only and is not intended for use in production systems.

Please note that content related to the previous px() equivocation function has been removed. A function, pxz3(), using projection, controlled erasure based equivocation, and recombination over related finite fields has been constructed that demonstrably creates many output to one input mappings. The formalization is a work in progress.

Contents

1 Introduction

1.1 The State of the Art

Modern quantum resilient cryptographic signature schemes are primarily based on structured lattice or hash based cryptography, each with a unique set of disadvantages. Lattice schemes, such as ML-DSA[2] leverage module lattices, which by virtue of their internal algebraic structure provide a potential avenue for quantum cryptanalysis, and variable sizes which are 5x the size of the signature primitives in use today. Hash based signature schemes such as SLH-DSA[1], to their credit, have virtually no exploitable algebraic structure and have small public keys, but produce relatively huge signatures requiring a substantial amount of computational resources to produce. The Falcon[11] signature algorithm, which has also been selected by NIST for standardization, is more communication efficient compared to ML-DSA, but relies on significantly more internal structure in the form of NTRU lattices, and remains challenging to implement in constant time due to reliance on floating point operations.

Signature solutions under consideration currently can generally be categorized as Lattice, Code based, Multivariate, Isogeny, and symmetric cryptography constructions with multiple variants remaining under consideration by NIST. Unfortunately, out of the 14 digital signature schemes selected for the NIST's second 'onramp' round of analysis, when considering communication cost alone, only SQISign[3] can be considered a reasonable replacement for current quantum weak signatures. Unfortunately, SQISign has the highest computational cost when compared to every other alternative, making widespread adoption fairly impractical. Additionally as isogeny based systems are fairly novel, and the predecessor SIDH was broken classically in a rather spectacular fashion[7], healthy skepticism is warranted. The other 13 candidates have similar or higher communications overhead compared to Falcon, and are based on largely untested computational hardness assumptions.

As the modern digital world has been built using digital signatures based on ECDSA and RSA, without a quantum resilient replacement with similar communications overhead, we are faced with significant challenges. We will have no choice but to redesign hardware, software, and protocols to accommodate these vastly increased communication costs, which will cost billions of dollars and take decades to migrate to. The current situation for digital signatures appears to be quite bleak. At this point, the probability that a provable quantum resilient digital signature solution that is both communication and computationally efficient will emerge is nearly inconceivable.

1.2 Secure and Efficient Signatures from the Non-Abelian Hidden Subgroup Problem and Information Theory

Non-commutative cryptography[10] has been an area of research since before the 1990s, and remains a credible basis for quantum resilient cryptographic systems. Previous attempts to construct secure non-commutative cryptographic systems were either communications inefficient compared to lattice schemes[8][4][6], or compromised by algebraic cryptanalysis. With most schemes having some form of structural weaknesses[9], the bulk of research inertia has since shifted to lattice or other well known problem groups with the perceived potential to achieve quantum resilience.

While there have been several attempts to leverage matrix groups for non-commutative cryptography[5], the chosen matrices were structured, lacked entropy, had exploitable subgroups, or exploitable linear relationships. These schemes relied on matrix equivalence or conjugacy problem hardness assumptions, which turned out to be less robust than anticipated. Our work uses a set of dense, full rank random matrices, forming a bilinear group. By using non-correlated matrices, we mitigate attacks leveraging structural relationships. With half of the matrix group randomized each operation, observed correlation only applies to a single matrix pair. This acts to effectively minimize exploitable relationship information. This bilinear matrix group structure form the building blocks of our instance of the non-Abelian Hidden Subgroup Problem (NAHSP) and is further strengthened by a novel equivocation function.

Non-commutative cryptographic schemes have received moderate attention within the cryptographic community, yet the Non-Abelian Hidden Subgroup Problem (NAHSP) remains comparatively under-explored. In contrast, the NAHSP is recognized as one of the most significant and challenging problems in quantum algorithmic research. For over two decades, intensive study of the NAHSP has not only failed to produce efficient quantum algorithms for general non-abelian groups, but has also yielded numerous negative results. Unlike problems that are merely conjectured to be quantum-resistant, the NAHSP has demonstrated resilience against decades of quantum attack attempts, establishing itself as a robust foundation for constructing quantum-resilient cryptographic systems.

In the context of this construction, the low dimensions $n = \{64, 128, 256\}$ and relatively small modulus sizes $q = \{257\}$ do not directly result in an intractable NAHSP instance. Normally, an adversary observing a small number of intermediate outputs could construct a solvable system of linear equations via modular arithmetic and Gaussian elimination. This attack vector is mitigated by equivocation operation. Per operation matrix randomization acts to uniformly distribute intermediate output elements of matrix calculations across a large ambient module space $\mathbb{Z}_n^q$. We leverage information-theoretic principles[13], equivocation and mutual entropy, to prove practical information theoretic security guarantees. While we do not claim the perfect secrecy of the one-time pad[12], we use information theoretic arguments to quantify the computational infeasibility of reconstructing the problem. We employ a novel equivocation function to partition the ambient space into equivalence classes, dispersing secret entropy across a class of indistinguishable valid pre-images.

This hybrid security model amplifies computational complexity with information-theoretic security. By leveraging uniform distribution and indistinguishability, we transform the computational problem of solving the NAHSP into an information-theoretic challenge. The challenge for an adversary shifts from using linear algebra to solve a well posed system of equations to identifying the correct indistinguishable elements needed to construct it.

In practical terms, this work challenges the prevailing notion that achieving quantum resilience necessitates a significant increase in communication overhead. Our construction integrates the structural advantages of non-abelian matrix groups with information-theoretic principles, offering an alternative direction for designing cryptographic systems that balance quantum resistance with real-world usability. The short signatures and public keys comparable to ECDSA and RSA may be especially useful in highly constrained network environments, where established lattice based schemes struggle to integrate.

This work is a preprint released to facilitate discussion and collaboration. Updates and refinements will follow as necessary.

2 High-Level Description

2.1 Core Representation

Our instance of the Non-Abelian Hidden Subgroup Problem (NAHSP) is defined within the context of a bilinear matrix group G. The group operation, matrix-vector multiplication, involves a public matrix A and a secret vector x, while the hidden subgroup N is defined through transformations induced by the secret matrix U. To amplify cryptographic hardness, the equivocation function $pxz()$ maps outputs to indistinguishable equivalence classes, obfuscating the relationships between A, U, x, and the observed results.

We can concisely represent key generation as:

$$t \equiv A \cdot x \mod q,$$

$$t' \equiv U \cdot t \mod q,$$

$$pk = t'' \equiv pxz(t') \mod q,$$

For signing as:

$$t \equiv B \cdot x \mod q,$$

$$t' \equiv U \cdot t \mod q,$$

$$sig = t'' \equiv pxz(t' \circ J(C1)) \mod q,$$

For validation as:

$$LHS \equiv B \cdot (pk \circ J(C2)) \mod q,$$

$$LHS' \equiv pxz(LHS \circ J(C1)) \mod q,$$

$$RHS \equiv A \cdot (sig \circ J(C2)) \mod q,$$

$$RHS' \equiv pxz(RHS) \mod q,$$

$$LHS' \stackrel{?}{=} RHS'$$

where:

A: Dense, full-rank, random public matrix used for key generation and verification.

B: Dense, full-rank, random public matrix used for signing and verification.

U: Dense, full-rank, random, uncorrelated private matrix.

x: Secret vector uniformly sampled from $\{1, ...q\}$ ensuring no zero entries.

t: Element in the right subgroup $H_{\text{right}} \subseteq G$, spanning the ambient space and serving as a hidden input to H_{left}.

$J()$: Secure hash function (e.g., SHA3/SHAKE) producing pseudorandom outputs in $\mathbb{Z}_q$.

t': Element in the hidden left subgroup $N \subseteq G$.

$pxz()$: Function mapping t' into equivalence classes, ensuring information theoretic infeasibility of t''s recovery.

t'': Obfuscated output element.

q: Prime modulus defining the finite field $\mathbb{Z}_q$.

fs: Fiat-Shamir heuristic in $\mathbb{Z}_q$, binding the public key, message, and randomness to the signature context.

r: Signature randomizer value in $\mathbb{Z}_q$, mitigating replay attacks and enhancing security against sEU-CMA.

pk: Public key element in $\mathbb{Z}_q$.

$C1$: Forgery constraint 1, an element in $\mathbb{Z}_q$, derived as a hash of intermediate context values related to $pk \cdot B$.

$C2$: Forgery constraint 2, an element in $\mathbb{Z}_q$, derived as a hash of pk, sig, and other context elements.

$\circ$: Element-wise multiplication operation.

$\cdot$: Matrix-vector product operation.

3 Problem Statement

The objective of this cryptographic scheme is to secure the hidden subgroup N, ensuring its structure remains concealed from adversaries. This is achieved by obfuscating the relationships between the public basis A, the hidden matrix U, and the secret vector x, while leveraging a lossy mapping function $pxz(\cdot)$ to induce equivalence classes. The scheme employs dual-layered security mechanisms: computational hardness from the Non-Abelian Hidden Subgroup Problem (NAHSP) and information-theoretic obfuscation from $pxz2(\cdot)$.

Adversary's Knowledge

The adversary has access to:

- The public matrix A, which spans the modular ambient space G_{ambient} and operates on H_{right},
- The final equivocated output t'', resulting from applying the lossy mapping $pxz2(\cdot)$ to elements of H_{left}.

Adversary's Limitations

The adversary does not have access to:

- The secret vector x, used in the transformation $t = A \cdot x \mod q$,
- The intermediate vector t, which resides within H_{right},
- The private matrix U, responsible for mapping elements from H_{right} to H_{left} and defining the hidden subgroup N.

Equivocation of the Mapping Function $pxz2(\cdot)$

The mapping function $pxz2(\cdot)$ is a lossy transformation that projects elements of the ambient group G_{ambient} into equivalence classes. This mapping introduces significant obfuscation, ensuring that subgroup membership cannot be determined feasibly without knowledge of the private components.

Properties of $pxz2(\cdot)$:

- $px : G_{\text{ambient}} \rightarrow$ Equivalence Classes, where each equivalence class contains indistinguishable pre-images.
- The function disrupts linear and algebraic relationships within the group, rendering coset structures unobservable.
- Without access to x or U, distinguishing subgroup members from non-members within equivalence classes is infeasible.

Impact of Equivocation: The lossy nature of $pxz2(\cdot)$ exponentially increases the adversary's search space by creating a many-to-one mapping:

- Pre-images of $pxz2(\cdot)$ form equivalence classes that mask coset relationships within H_{left},
- The adversary must contend with an exponentially large number of indistinguishable elements, effectively reducing any observed output to noise.
- By the *Data Processing Inequality (DPI)*, a result derived from Shannon's foundational work in information theory (Theorem of Noisy Channels), the mutual entropy between the input and output of $pxz2(\cdot)$ is provably reduced through this lossy mapping. The surjectivity of $pxz2(\cdot)$ increases the likelihood that the output's entropy is maximized relative to the adversary's view, rendering it statistically indistinguishable from random noise and enhancing equivocation.

Chaining Mechanism

The chaining mechanism enhances security by linking independent problem instances through intermediate outputs. Each stage introduces unique secrets and transformations, ensuring that the overall system is resilient against adversarial attacks. To give an example in the signature context with six chained instances $k = 6$:

For $k = 0$:

$$t_0 \equiv B_0 \cdot x_0 \mod q,$$
$$t'_0 \equiv U_0 \cdot t_0 \mod q,$$
$$t''_0 \equiv pxz2(t'_0 \circ J(C1_0)) \mod q,$$

For $k = 1$ to 5:

$$t_k \equiv B_k \cdot (x_k \circ t''_{k-1}) \mod q,$$
$$t'_k \equiv U_k \cdot t_k \mod q,$$
$$t''_k \equiv pxz2(t'_k \circ J(C1_k)) \mod q,$$

Where the final output $sig = t''_5$.

Mechanism:

- Intermediate outputs t''_k from one stage are passed as hidden inputs to the next stage,
- Each stage employs independent secrets x_k, matrices U_k, and public matrices A_k, ensuring randomness and unlinkability.

Security Benefits:

- Error Propagation: Any errors or approximations in recovering one stage amplify across subsequent stages, compounding the adversary's difficulty.
- Independence of Stages: Knowledge of secrets from one stage does not simplify reconstruction of secrets from subsequent stages due to the introduction of fresh randomness and transformations.
- Unlinkability: Intermediate values t_k and t'_k remain hidden, ensuring that adversaries cannot correlate outputs across stages.
- Hardness Amplification: Solving one instance of the chained system yields no useful information for subsequent stages. The adversary must solve all instances simultaneously, which exponentially increases the overall complexity of the problem.

Verification and Forgery Mitigation

The verification process ensures the integrity of the transformations applied during signing and key generation, validating that the observed signature σ corresponds to the public key pk and the private components x and U, without revealing these private components. By leveraging the mapping function $pxz2(\cdot)$, contextual hash constraints, and entropy checks, the scheme mitigates forgery attempts while maintaining soundness and completeness.

Verification Equation: The verification process compares two transformed outputs derived from the public key pk and the signature σ, iteratively refining them under contextual constraints $C1$ and $C2$:

- $C1_k$: Represents the intermediate value derived during signing and verification, computed as $J(pk \cdot B_k)^3$, ensuring consistency between signing and verification.
- $C2_k$: A hash of pk, σ, FS (Fiat-Shamir Heuristic), and r (message randomizer), binding the signature to its context and mitigating signature malleability.

The verification equation is computed as follows:

$$LHS_0 \leftarrow pk, \quad RHS_0 \leftarrow \sigma$$

For $k = 0$ to $k - 2$:

$$LHS_{k+1} = pxz3(B_k \cdot (LHS_k \circ J(C2_k)) \mod q), \quad RHS_{k+1} = pxz3(A_k \cdot (RHS_k \circ J(C1_k)) \mod q).$$

For the final stage ($k = k - 1$):

$$LHS_k = B_k \cdot (LHS_{k-1} \circ J(C2_k)) \mod q, \quad RHS_k = A_k \cdot (RHS_{k-1} \circ J(C1_k)) \mod q.$$

A signature σ is valid if and only if:

$$LHS_k \overset{?}{=} RHS_k.$$

Key Components:

- Public Matrices A and B: Define the observable transformations applied to the public key and signature during verification.
- Secure Hash Function $J()$: Produces contextual constraints $C1$ and $C2$, binding the signature and public key to the specific signing context.
- Mapping Function $pxz3(\cdot)$: Compresses equivalence classes to ensure subgroup membership remains indistinguishable, preventing adversarial reconstruction of x or U.
- Observed Entropy Constraint: During both signature generation and validation, the observed entropy and randomness of signature candidates are checked to ensure they statistically represent approximately 1/10 of the combinatorial possibilities. This constraint aligns with information-theoretic principles, ensuring that signatures exhibit near-randomness and resist predictability.

Forgery Mitigation: Forgery is mitigated through the interaction of several mechanisms:

- Contextual Hash Constraints: The hash constraints $C1$ and $C2$ bind the signature and public key to specific contextual values, ensuring that signatures cannot be reused or manipulated across different contexts.
- Lossy Mapping Function and Probability of Forgery: The fixed lossy mapping function $pxz3(\cdot)$ acts to induce a lower level of equivocation than $pxz2(\cdot)$, making it computationally infeasible for adversaries to generate valid signatures without access to the private keys. The probability of a successful forgery passing each level k is determined by the ratio of equivalence class members m_k to the total number of equivalence classes n_k. For each level, the adversary must generate a value that maps to the correct equivalence class under $pxz3(\cdot)$, resulting in a success probability of approximately $\frac{m_k}{n_k}$. Across K levels, the cumulative probability of forging a valid signature is given by:

  $$P_{\text{forgery}} \sim \prod_{k=0}^{K-1} \frac{m_k}{n_k}.$$

  This product reflects the compounding difficulty of forging a signature, as the adversary must satisfy all constraints simultaneously. The carefully chosen ratio of equivalence class members to equivalence class numbers ensures that the probability of a successful forgery remains negligibly small.

- Fixed and Randomized Public Matrices: The public matrix A is fixed and defines the ambient group structure, while the signing matrix B is randomized and tied to the message. This dynamic prevents adversaries from correlating multiple signatures to infer structural relationships or exploit reuse.
- Observed Entropy Constraint: The additional proposed observed entropy constraint ensures that edge-case scenarios are effectively mitigated. Signature candidates are required to statistically adhere to near-randomness, aligning with approximately $1/10th$ of the combinatorial possibilities. This increases the difficulty of identifying predictable or exploitable patterns, enhancing resilience to forgery.

Soundness and Completeness:

- Completeness: Any valid signature σ, generated using the correct set of private components x and U, satisfies the verification equation.
- Soundness: Any invalid signature σ', generated without access to the private components, fails verification. This failure arises because σ' maps to incorrect equivalence classes under $pxz3(\cdot)$, and fails entropy checks for statistical validity.

Related proofs of soundness, completeness, equivocation, and equivalence class ratio impact will be presented as part of the formal proof of existential unforgeability under chosen message attacks (EUF-CMA) in a subsequent section.

Adversary's Tasks: Key Recovery vs. Forgery

The adversary's objectives can be categorized as follows:

- Key Recovery: Reconstructing the hidden subgroup N by recovering U and x:
  - This requires solving the NAHSP, an infeasible task due to the equivocation induced by $pxz2(\cdot)$ and the computational hardness of the problem.
- Forgery: Generating a valid signature σ' without access to the private keys:
  - This requires reversing branching layers of $pxz3(\cdot)$ to identify valid subgroup elements, which is infeasible due to the lossy nature of the mapping and the randomness introduced at each stage.

Summary

The scheme achieves robust security by:

- Obfuscating subgroup structure through the lossy mapping $pxz2(\cdot)$,
- Amplifying adversarial difficulty with the chaining mechanism, ensuring that errors propagate across stages,
- Maintaining soundness and completeness in the verification process, ensuring only valid signatures satisfy the verification equation,
- Combining computational hardness from the NAHSP with information-theoretic obfuscation from $pxz2(\cdot)$, ensuring resilience against both classical and quantum adversaries.

4 Preliminary Results and Contributions

This work introduces a novel digital signature scheme that incorporates both **information-theoretic security** and **computational hardness**, explicitly tied to the Non-Abelian Hidden Subgroup Problem (NAHSP). While the results presented are preliminary, they suggest a promising approach to balancing efficiency, security, and practicality in post-quantum cryptography. The key contributions of this work include:

1. Fiat-Shamir Transformation with Contextual Binding: Leveraging the Fiat-Shamir transformation to securely bind the public key, message, and randomness, generating a challenge seed that derives a unique set of challenge bases per signature. This approach reinforces security and ensures contextual linkage between the signature and the corresponding public key.

2. Chaining Mechanism for Amplified Hardness: Introducing a chaining mechanism that combines independent instances of the matrix-based NAHSP problem. Each stage introduces fresh randomness and transformations, compounding adversarial complexity and amplifying computational difficulty with every additional stage.

3. Verification through Structured Basis Transformations: Designing signature verification as a proof of consistency through structured basis transformations. This approach ensures that transformations applied to the public key and signature align under obfuscated subgroup relationships, preserving algebraic correctness while preventing exploitation of subgroup structures.

4. Information-Theoretic Security via High-Entropy Mapping Functions: Introducing a non-linear, many-to-one mapping function $pxz2(\cdot)$ that compresses the NAHSP output in ambient space G_{ambient} into equivalence classes. The inherent high entropy of $pxz2(\cdot)$'s outputs enforces:

   - Computational indistinguishability of equivalence class elements without access to private keys x_k and U_k,
   - Explicit rejection of low-entropy forgeries during verification, adding an entropy-based security layer that complements computational hardness.

5. Novel Matrix-Based Cryptographic Framework: Developing a cryptographic platform based on unrelated public/private matrix groups. The independence of public matrices A, B, and private matrices U, x prevents structural exploitation, supporting efficient signature generation and verification.

6. First Practical Hybrid Information Theoretic and NAHSP-Based Construction: Constructing what we believe to be the first practical digital signature scheme explicitly based on the NAHSP, using non-commutative matrix groups and leveraging both information theoretic functions and chaining mechanisms to ensure robust security.

Context and Preliminary Results: The proposed scheme, though unrefined, demonstrates the potential of non-commutative cryptography to address critical challenges in quantum resilience. While further validation, cryptanalysis, and exploration of parameter optimization are necessary, the approach offers:

- Communication Efficiency: Preliminary parameters suggest competitive communication costs compared to RSA signatures and a significant reduction compared to most lattice-based schemes.

- Dual-Layered Security: By combining information-theoretic indistinguishability with computational hardness rooted in the NAHSP, the scheme provides a layered defense against both classical and quantum adversaries.

- Feasibility and Scalability: The use of independent, unrelated public/private matrix groups provides a scalable and tunable platform for balancing security and efficiency, with conservative parameter choices supporting incremental improvements over time.

Careful Optimism: While matrix group-based schemes have been explored in the past, this work introduces novel techniques that warrant renewed investigation of non-commutative cryptography. The preliminary results presented here are encouraging but must be rigorously validated by the broader cryptographic community. Future work will focus on parameter tuning, formal proof refinement, performance improvements, and independent verification to solidify the scheme's practical viability and theoretical soundness.

Complexity Analysis The mapping function $pxz2(\cdot)$ partitions the group G into equivalence classes, obfuscating the subgroup structure of N and increasing the adversary's difficulty in distinguishing elements. Unlike a group homomorphism, $pxz2(\cdot)$ does not preserve group operations but ensures computational indistinguishability of elements within the same equivalence class. This indistinguishability amplifies the complexity of solving the problem by significantly increasing the effective solution space.

Impact of $pxz2(\cdot)$: The hardness of the problem is tied directly to the pre-image count of $pxz2(\cdot)$, which determines the size of equivalence classes and the adversarial search space. Specifically:

- The adversary must navigate all elements in $px^{-1}(g')$ for a given equivalence class g', where $px^{-1}(g')$ contains all pre-images indistinguishable under $pxz2(\cdot)$.

- The size of $px^{-1}(g')$ is determined by the partitioning of the ambient group G_{ambient} into equivalence classes via the mapping $pxz2(\cdot)$. The hidden subgroup N and the transformations induced by x and U influence the structure of these partitions, but the pre-image membership size fundamentally scales with the size of G_{ambient} and configuration of $pxz2(\cdot)$.

- The indistinguishability within equivalence classes ensures that structural relationships between elements of N and G are practically obscured, limiting adversarial insights.

Chaining Mechanism and Amplified Complexity: The chaining mechanism compounds complexity by propagating errors and introducing additional randomness at each stage, requiring the adversary to solve multiple independent instances of the obfuscated problem. In a single instance, the complexity of solving the problem scales with the size of equivalence classes induced by $pxz2(\cdot)$. For k chained instances, the total complexity is amplified as:

$$O(|px^{-1}(g')|^k),$$

where $|px^{-1}(g')|$ is the size of the pre-image set for a single equivalence class. This reflects:

- The exponential growth of the adversarial search space due to chained instances, requiring reconstruction of intermediate outputs to solve subsequent stages.
- The cascading effect of errors, where small inaccuracies in earlier stages propagate, significantly increasing the difficulty of reconstructing the entire system.

Security and Complexity Relationship: The indistinguishability introduced by $pxz2(\cdot)$ ensures that adversaries cannot efficiently distinguish elements of N within equivalence classes or between stages of the chaining mechanism. By explicitly tying the complexity to the pre-image count $|pxz2^{-1}(g')|$, the scheme achieves:

- Scalable Hardness: The size of equivalence classes and the number of chained instances can be tuned to balance efficiency and security.
- Resistance to Structural Attacks: The obfuscation introduced by $pxz2(\cdot)$ disrupts structural relationships, ensuring that subgroup recovery requires infeasible computational resources.
- Cascading Complexity: The chaining mechanism amplifies the adversarial challenge, requiring reconstruction of intermediate outputs across multiple stages, with errors compounding exponentially.

Practical Observations: This work does not claim proven hardness for the NAHSP in the general case but leverages its empirical resistance to both classical and quantum attacks. The inclusion of $pxz2(\cdot)$ and chaining mechanisms provides additional layers of security, making the scheme robust under practical cryptographic assumptions while maintaining tunable efficiency for real-world applications.

Communication Cost: Perhaps the most relevant result of this work is leveraging information theoretic security to achieve practical signature and public key sizes.

Level	n	**PK (bytes)**	**Sig (bytes)**	k
I	64	80	96	8
III	128	152	176	6
V	256	288	320	4

Table 1: Public Key, Signature Sizes, and Chain Instances Across Levels

4.1 Structure of Remainder of Paper

- Formal mapping of our construct to the Non-Abelian Hidden Subgroup Problem
- Analysis of the information theoretic properties of $pxz2(\cdot)$ in relation to NAHSP
- Proof of Verification Constancy

- Security Notes and Attack Models
- Proof of IND-CPA hardness
- Proof of sEU-CMA security
- Algorithms and Implementation
- Performance Comparison
- Future Work and Conclusion

5 Notes

Throughout this paper, we give both abstract parameters and concrete example formulas. If specific values are used, they are based on the level III instance, as thus far it has received the bulk of analysis.

Additionally, if you are familiar with the Number Theoretic Function, we recommend you that you do not assume too much. Some readers have gotten hung up on the assumption that NTTs are by definition invertible and can only be used to construct invertible functions.

6 Formal Reduction to the Non-Abelian Hidden Subgroup Problem (NAHSP)

For this mapping, it is important to note that every element has already been forward transformed to its NTT representation, and essentially all operations are performed in the base 'NTT domain'. Generally this base domain is defined by the NTT configured using $q = 257; \omega = 3$. The NTT imparts isomorphic structure to elements.

6.1 Group Structure and Properties

Ambient Group $G_{\textbf{ambient}}$: The ambient group G_{ambient} is defined as the modular output space where elements of the hidden subgroup reside. Formally, it is represented as:

$$G_{\text{ambient}} = \mathrm{GL}(n, \mathbb{Z}_q) \rtimes \mathbb{Z}_q^n,$$

where the group operation for elements (A, x) and (U, y) in G_{ambient} is given by:

$$(A, x) \cdot (U, y) = (AU, Ay + x) \mod q.$$

This group is non-abelian for $n \geq 2$ due to the non-commutative nature of matrix multiplication in $\mathrm{GL}(n, \mathbb{Z}_q)$. The semi-direct product structure allows for incorporating both linear transformations (via matrices) and translations (via vectors), which is essential for modeling the transformations in the NAHSP scheme.

Group G: The composite group G is constructed as a semi-direct product of two cyclic subgroups derived from specific matrices involved in the transformation equations:

$$G = H_{\text{left}} \rtimes H_{\text{right}},$$

where:

- $H_{\text{left}} = \langle U \rangle$, the cyclic subgroup generated by the matrix U,
- $H_{\text{right}} = \langle A \rangle$, the cyclic subgroup generated by the matrix A.

Here, A and U are invertible matrices in $\mathrm{GL}(n, \mathbb{Z}_q)$, and their powers generate the respective cyclic subgroups. The semi-direct product ensures that G is non-abelian, provided that the action of H_{right} on H_{left} is non-trivial.

Group Operation in G: The group operation within G combines elements from H_{left} and H_{right} as follows. For any $h_L \in H_{\text{left}}$ and $h_R \in H_{\text{right}}$:

$$h_R \cdot h_L \cdot h_R^{-1} = \phi_{h_R}(h_L),$$

where ϕ_{h_R} is an automorphism of H_{left} induced by conjugation by h_R. Specifically, if $h_R = A^k$ and $h_L = U^m$, then:

$$A^k U^m A^{-k} = (A^k U A^{-k})^m.$$

This relation encapsulates how elements of H_{right} act on H_{left}, ensuring the non-abelian nature of G.

Incorporating Transformation Equations: The transformation equations central to the NAHSP scheme are:

$$t = Ax \mod q, \quad t' = Ut \mod q.$$

These can be interpreted within the group structure as follows:

- First Transformation: Applying A to a vector x corresponds to the action of $A \in H_{\text{right}}$ on $x \in \mathbb{Z}_q^n$.
- Second Transformation: Applying U to the result t corresponds to the action of $U \in H_{\text{left}}$ on t.

Thus, the sequential application of A and U reflects the group operation in G_{ambient}, where each transformation is represented by elements from H_{right} and H_{left}, respectively.

Properties Ensuring Group Validity: To confirm that G is indeed a group under the defined operation when performed in the NTT domain, we verify the group axioms:

- Closure: For any $h_R = A^k \in H_{\text{right}}$ and $h_L = U^m \in H_{\text{left}}$, their product $h_R h_L = A^k U^m$ is also in G, as $A^k \in H_{\text{right}}$ and $U^m \in H_{\text{left}}$.
- Associativity: Follows from the associativity of matrix multiplication and vector addition in $\mathrm{GL}(n, \mathbb{Z}_q) \rtimes \mathbb{Z}_q^n$.
- Identity Element: $(I, 0)$, where I is the identity matrix in $\mathrm{GL}(n, \mathbb{Z}_q)$ and 0 is the zero vector in $\mathbb{Z}_q^n$, serves as the identity element.
- Inverse Element: For each $(A^k, U^m) \in G$, the inverse is (A^{-k}, U^{-m}), ensuring that $(A^k, U^m) \cdot (A^{-k}, U^{-m}) = (I, 0)$.

Non-Abelian Nature: The group G is non-abelian provided that A and U do not commute, i.e., $AU \neq UA$. This non-commutativity is crucial for the complexity inherent in the NAHSP, as it relies on the hidden subgroup being non-abelian for quantum resilience.

6.2 Application to NAHSP Scheme

In the context of the NAHSP scheme, the transformations $t = Ax \mod q$ and $t' = Ut \mod q$ are modeled by the group operations within G. Specifically:

$$\begin{aligned} t = Ax \mod q \quad &\text{corresponds to} \quad (A, 0) \cdot (I, x) = (A, Ax \cdot 0 + x) = (A, Ax), \\ t' = Ut \mod q \quad &\text{corresponds to} \quad (U, 0) \cdot (A, x) = (UA, Ux + 0) = (UA, Ux). \end{aligned}$$

This sequential application of group elements $(A, 0)$ and $(U, 0)$ encapsulates the transformations defined in the NAHSP scheme, aligning the algebraic structure with the operational procedures of the problem.

6.3 Summary

Ambient Group G_{ambient}: Combines linear transformations and translations in the NTT domain through a semi-direct product, enabling the representation of both A and U within a unified group framework.

Group G: A non-abelian subgroup constructed from cyclic subgroups generated by A and U, facilitating the structured application of transformations central to the NAHSP scheme.

Group Operations: Accurately model the transformation sequences $t = Ax \mod q$ and $t' = Ut \mod q$, ensuring that the group axioms are satisfied and the non-abelian properties are maintained.

Hidden Subgroup N: The hidden subgroup N is defined as $N = H_{\text{left}}$. As established in Lemma 2, N is a normal subgroup of G, ensuring that:

$$gNg^{-1} \subseteq N, \quad \forall g \in G.$$

This normality guarantees that G can be partitioned into disjoint cosets of N:

$$G = \bigcup_i g_i N, \quad g_i N \cap g_j N = \emptyset \text{ for } i \neq j.$$

6.4 Formal Definition of NAHSP

Definition 1 (Non-Abelian Hidden Subgroup Problem (NAHSP))**.** *The* ***Non-Abelian Hidden Subgroup Problem*** *(NAHSP) is defined as follows:*

- ***Input:*** *A finite non-abelian group G and a function $f : G \to S$, where S is the set of left cosets of a hidden subgroup $H \subseteq G$. The function f satisfies:*

$$f(g) = f(g') \iff gH = g'H.$$

- ***Output:*** *Determine the hidden subgroup H.*

6.5 Oracle Construction and Reduction to NAHSP

Oracle Function f: Define the oracle function $f : G \to S$ as:

$$f(g) = gN,$$

where gN is the left coset of N containing g. The function f satisfies the equivalence relation:

$$f(g) = f(g') \iff gN = g'N.$$

Thus, f is constant on cosets of N and distinct across cosets, fulfilling the requirements of the NAHSP.

Correspondence with NAHSP:

- **Group G:** The group $G = H_{\text{left}} \rtimes H_{\text{right}}$ serves as the non-abelian group in the NAHSP framework.
- **Hidden Subgroup H:** The hidden subgroup H in NAHSP corresponds to $N = H_{\text{left}}$ in our construction.
- **Oracle Function f:** The oracle function $f(g) = gN$ encodes the coset structure of N, aligning with the oracle requirements of NAHSP.

6.6 Hardness of Subgroup Recovery

- **Non-Abelian Structure:** The semi-direct product $G = H_{\text{left}} \rtimes H_{\text{right}}$ is inherently non-abelian due to the automorphism action of H_{right} on H_{left}. This non-abelian nature prohibits the direct application of abelian techniques, such as Fourier analysis, which are pivotal in efficiently solving the Hidden Subgroup Problem (HSP) in abelian groups.
- **Classical Complexity:** Classical algorithms lack the necessary tools to exploit the group structure effectively. They would be compelled to perform exhaustive brute-force enumeration over the cosets of N, a task rendered computationally infeasible by the exponential size of G. Moreover, the intertwined structures of H_{left} and H_{right} offer no combinatorial shortcuts for efficient subgroup identification.
- **Quantum Complexity:** Quantum algorithms, particularly those utilizing the Quantum Fourier Transform (QFT), falter in non-abelian settings like G. The automorphism action of H_{right} on H_{left} disrupts the coherence and periodicity necessary for QFT-based techniques to identify subgroup structures efficiently. Consequently, these quantum approaches do not yield a polynomial-time solution for NAHSP in such non-abelian groups.

Conclusion

Recovering the hidden subgroup $N = H_{\text{left}}$ in the group $G = H_{\text{left}} \rtimes H_{\text{right}}$ satisfies the definition of the Non-Abelian Hidden Subgroup Problem (NAHSP). The non-abelian structure of G, combined with the automorphism action of H_{right} on H_{left}, ensures that this problem is computationally infeasible under both classical and quantum adversarial models. Thus, the cryptographic hardness of the NAHSP is directly inherited by the problem of recovering N in G. □

6.7 Reduction

Adversarial Setup. Let $\mathcal{A}$ be an adversary attempting to recover the hidden subgroup $N = H_{\text{left}}$ from the group $G = H_{\text{left}} \rtimes H_{\text{right}}$. The adversary interacts with an oracle function $f : G \to S$, where S is the set of left cosets of N in G. The function f is defined as:

$$f(g) = gN,$$

where gN is the coset of N containing g. The function f satisfies the equivalence relation:

$$f(g) = f(g') \iff gN = g'N.$$

The adversary's goal is to identify N given oracle access to f.

Definition of Security. The adversary's advantage $\text{Adv}_{\mathcal{A}}$ in recovering N is defined as:

$$\text{Adv}_{\mathcal{A}} = \Pr[\mathcal{A}(f) = N] - \Pr[\mathcal{A}_{\text{random}}(f) = N],$$

where $\mathcal{A}_{\text{random}}$ is a baseline adversary that outputs a random subgroup N' chosen uniformly at random from the set of all possible subgroups of G. The probabilities are taken over the random choice of N and any randomness inherent in the adversaries.

Reduction to the Non-Abelian Hidden Subgroup Problem. Assume $\mathcal{A}$ is an adversary that can recover the hidden subgroup $N = H_{\text{left}}$ with advantage ϵ. We construct a reduction $\mathcal{R}$ that uses $\mathcal{A}$ to solve the Non-Abelian Hidden Subgroup Problem (NAHSP) as follows:

1. **Input to $\mathcal{R}$:** The group $G = H_{\text{left}} \rtimes H_{\text{right}}$ and oracle function $f : G \to S$ defined by $f(g) = gN$, where $N = H_{\text{left}}$.

2. **Reduction Steps:**

   (a) $\mathcal{R}$ provides $\mathcal{A}$ with oracle access to f.

   (b) $\mathcal{A}$ outputs a candidate subgroup N'.

   (c) $\mathcal{R}$ verifies whether N' is a valid hidden subgroup by checking:

   $$\forall g, g' \in G, \quad g^{-1}g' \in N' \iff f(g) = f(g').$$

   This ensures that N' correctly defines the coset structure as per the oracle f.

3. **Output of $\mathcal{R}$:** If verification succeeds, $\mathcal{R}$ outputs N' as the solution to NAHSP. Otherwise, $\mathcal{R}$ outputs failure.

Analysis of the Reduction.

- **Correctness:** If $\mathcal{A}$ successfully identifies N, then $\mathcal{R}$ correctly solves NAHSP by outputting $N' = N$. The verification step ensures that N' uniquely satisfies the coset equivalence relation defined by f, thereby guaranteeing the correctness of the solution.

- **Efficiency:** The reduction $\mathcal{R}$ invokes $\mathcal{A}$ once and performs polynomial-time group operations for verification. Therefore, the computational overhead of $\mathcal{R}$ is polynomial in the size of G and bounded by the runtime of $\mathcal{A}$.

- **Adversarial Advantage:** Suppose $\mathcal{A}$ has a non-negligible advantage ϵ in recovering N. Then, $\mathcal{R}$ achieves the same advantage in solving NAHSP:

  $$\text{Adv}_{\mathcal{R}} = \text{Adv}_{\mathcal{A}} = \epsilon.$$

  This implies that any adversary $\mathcal{A}$ capable of recovering N with advantage ϵ enables $\mathcal{R}$ to solve NAHSP with the same advantage.

Hardness of Subgroup Recovery.

- **Classical Adversaries:** Classical algorithms would need to enumerate cosets of N, which is computationally infeasible due to the exponential size of G. Additionally, the non-abelian structure of G lacks the necessary algebraic properties that allow for efficient subgroup identification, preventing the use of techniques such as brute-force search or combinatorial optimizations.

- **Quantum Adversaries:** Quantum algorithms, including those leveraging the Quantum Fourier Transform (QFT), struggle with G's non-abelian structure. The automorphism action of H_{right} on H_{left} disrupts the periodicity and coherence essential for QFT-based subgroup recovery. As a result, these quantum techniques fail to efficiently exploit the hidden subgroup structure in G, ensuring resistance against known quantum attacks.

Conclusion. This reduction demonstrates that recovering the hidden subgroup $N = H_{\text{left}}$ in $G = H_{\text{left}} \rtimes H_{\text{right}}$ is at least as hard as solving the Non-Abelian Hidden Subgroup Problem (NAHSP). The intractability of NAHSP under both classical and quantum adversarial models ensures the cryptographic security of the proposed system. □

6.8 Equivocation Function $pxz2(\cdot)$

This section is being re-formalized and will be available as soon as possible.

7 Proof of Consistency as Verification Under Homomorphic Transformations

Lemma 1. *The verification equation $LHS' = RHS'$ holds if and only if the signature σ is generated using the corresponding private keys and the specified public key $pk\|fs$, with high probability.*

Proof. Let the key generation, signing, and verification functions be defined as follows:

1. Key Generation

$$t = A \cdot x \mod q, \quad t' = U \cdot t \mod q, \quad pk = pxz2(t') \mod q,$$

where:

- A is a public matrix,
- x is the secret key,
- U is a private matrix,
- $pxz2$ is the full mapping function.

2. Signature Generation

$$\sigma = pxz2\big(U \cdot (J(C1) \circ t)\big) \mod q,$$

where:

- J is a hash function (e.g., SHAKE),
- $C1$ is constraint1, derived from $pk \cdot B$ intermediates after cubing and hashing.
- $\circ$ denotes a Hadamard product.

3. Verification Function

$$\begin{aligned} \text{LHS} &= B \cdot (pk \circ J(C2) \mod q, \\ \text{LHS}' &= pxz3\big(\text{LHS} \circ J(C1)\big) \mod q, \\ \text{RHS} &= A \cdot (\sigma \circ J(C2)) \mod q, \\ \text{RHS}' &= pxz3(\text{RHS}) \mod q. \end{aligned}$$

Step 1: Valid Signature Consistency

- Substitute the signature generation equation into RHS:

$$\text{RHS} = A \cdot \big(pxz3(U \cdot (J(C1) \circ t)) \circ J(C2)\big) \mod q.$$

- Using the properties of $pxz3$ and $t' = U \cdot t$, it follows that:

$$pxz3\big(U \cdot (J(C1) \circ t)\big) = pxz3(t') \mod q,$$

where t' satisfies the public key equation $pk = pxz3(t') \mod q$. - Therefore, the transformations applied during signing and verification align, yielding:

$$\text{RHS}' = pxz3(\text{RHS}) = pk \mod q.$$

Step 2: Validating LHS$'$

- Substitute pk into LHS:

$$\text{LHS} = B \cdot (pk \circ J(C2) \mod q.$$

- Apply the transformation $pxz3$:

$$\text{LHS}' = pxz3(\text{LHS} \circ J(C1) \mod q.$$

- Since the signature σ was generated using the correct private key, the transformations $J(C2)$ compensate for modular inconsistencies, ensuring:

$$\text{LHS}' = pk \mod q.$$

Step 3: Equivalence of LHS$'$ and RHS$'$

- Both LHS$'$ and RHS$'$ reduce to $pk \mod q$, implying:

$$\text{LHS}' = \text{RHS}' \iff \sigma \text{ was generated using the correct private key.}$$

Step 4: Probabilistic Argument for Invalid Signature

- For an invalid σ, the transformations in LHS and RHS will not align. To quantify this: - The output of $J(pk\|fs)$ is uniformly distributed over its range. - Each σ candidate not generated with the correct private key maps to a random equivalence class under $pxz3$, with negligible probability of aligning with LHS. - The adversary must guess both:

- σ, which depends on the secret key x and the private matrix U,
- The hash $J(pk\|fs)$, which is computationally infeasible due to the pre-image resistance of J.

- The success probability of forging σ without knowledge of x is bounded by:

$$P_{\text{success}} \leq \frac{1}{q^n},$$

where q^n is the size of the search space for σ. This represents an information-theoretic lower bound on the success probability.

Step 5: Contrapositive

- For an invalid σ, the mismatch between LHS$'$ and RHS$'$ occurs due to inconsistencies in equivalence class mapping, leading to:

$$\text{LHS}' \neq \text{RHS}'.$$

Conclusion

The verification equation LHS$'$ = RHS$'$ holds if and only if the signature σ is generated using the valid private key x, the private matrix U, and the specified public key and basis B constraint $C1$. The probabilistic argument establishes that forging a valid σ without knowledge of the private key is computationally infeasible with high probability. □

8 Implementation Details

8.1 Matrix Generation Using Diverse Cryptographically Secure PRNGs

To ensure cryptographic security and reproducibility, the public and private matrices in our construction should be generated deterministically using distinct cryptographically secure pseudorandom number generators (CSPRNGs). These are recommendations for high security, and certain implementations may prefer alternate functions.

8.1.1 Public Matrix Generation

The public matrices A used to generate the subgroup H_{right} are derived using AES-DRBG, per NIST-approved DRBG specifications. Each matrix $A \in \mathbb{Z}_q^{n \times n}$ is constructed as follows:

1. Input: A 256-bit public seed Seed_A, which may be application-specific or predefined.
2. Generation: Use AES-DRBG in CTR mode to generate n^2 entries.
3. Mapping: Map each entry modulo q to produce a dense, full-rank matrix A.
4. Validation: Optionally verify A's rank to ensure it is full rank.

This deterministic process is efficient, ensures reproducibility, and eliminates reliance on weak randomness.

8.1.2 Private Matrix Generation

The private matrices U, which define the subgroup H_{left}, are generated using SHA-512, SHA3-512, or SHAKE-256:

1. Pre-Input: Optionally use a private 256-bit (or larger) value to key the hash function.
2. Input: A 256-bit private seed Seed_U, derived from an entropy source or securely exchanged during key generation.
3. Hashing: Apply the chosen hash function to Seed_U to produce n^2 pseudorandom outputs.
4. Mapping: Map these outputs modulo q to construct U, ensuring full rank and density.
5. Validation: Optionally verify U's rank to confirm full rank.

8.1.3 Security Implications

Using AES-DRBG for public matrices and SHA-512/SHA3/SHAKE for private matrices ensures high entropy, cryptographic security, compliance with NIST standards, and diversity in matrix generation. These methods eliminate correlations between A and U, ensuring the subgroup structures H_{right} and H_{left} align with the theoretical reductions to NAHSP. Deterministic generation guarantees that the matrices are free from vulnerabilities introduced by weak or biased randomness. Furthermore, ensuring full rank for both matrices preserves the cryptographic strength of the construction.

8.2 Algorithm Details

8.3 Utility Algorithms

Note that we don't specify the specific pseudo-random algorithm used to expand the *seed* value, as this function is designed to be modular. In our reference instance we use AES256-DRBG, but other PRNG constructions are certainly supported.

Algorithm 1 Sample(*seed*)

Generates a matrix for the with non-zero elements.

Require: Dimensions K, N, prime modulus $Q1$, root $R1$, and seed `seed`.

```
Initialize $A[K][N][N]$ as an empty matrix.
for mat $= 0$ to $K - 1$ do
    rows_written $\leftarrow 0$.
    while rows_written $< N$ do
        Generate pseudo-random buffer buff using seed.
        for $y = 0$ to $N - 1$ do
            Extract trial_vec from buff.
            Apply transformation NTT(trial_vec, $Q1, R1$).
            if trial_vec contains no zero elements then
                Store trial_vec in $A$[mat][rows_written].
                rows_written $\leftarrow$ rows_written $+ 1$.
                if rows_written $= N$ then
                    break inner loop.
                end if
            end if
        end for
    end while
end for
return $A$.
```

Algorithm 2 genC2($elm1, elm2, m, SIG\_fs, SIG\_r$)

Generates a constraint element v by hashing inputs and reducing modulo $Q1$.

```
Require: Elements elm1, elm2 of size $N$, public variables $m$, SIG_fs, SIG_r of size
    SEED_SIZE, and prime modulus $Q1$.
Ensure: Element $v[N]$ with non-zero elements.
Initialize SHAKE256 context: mdctx.
if mdctx initialization fails then
    Throw error and terminate.
end if
Begin SHAKE256 hashing process.
Update hash with elm1, elm2, $m$, SIG_fs, and SIG_r.
Finalize hash to produce hash_output[$N \times$ sizeof(int32_t)].
for $i = 0$ to $N - 1$ do
    Extract $val$ from hash_output[$i$].
    Compute $v[i] \leftarrow \text{abs}(val) \mod Q1$.
    if $v[i] = 0$ then
        Set $v[i] \leftarrow 1$ to ensure non-zero component.
    end if
end for
Free SHAKE256 context: mdctx.
return $v$.
```

Algorithm 3 genC1(pk, SIG_MATRIX)

Generates constraining element set $C1$ by computing a cubed and hashed version of $pk \cdot B$.

Require: Element $pk[N]$, signature matrix SIG_MATRIX$[K][N][N]$, and prime modulus $Q1$.

Ensure: Matrix $C1[K][N]$ with processed values.

```
Initialize vector $LHS \leftarrow pk$.
Initialize $result[N] \leftarrow 0$, $LHS[N] \leftarrow pk$.
for mat = 0 to $K - 1$ do
    Reset $result[N] \leftarrow 0$.
    result $\leftarrow$ MatrixVectorProduct($\mathrm{SIG\_MATRIX[mat], LHS}, Q1$).
    LHS $\leftarrow$ result.
    Compute $LHS \leftarrow LHS^3 \mod Q1$.
    Initialize SHAKE256 context: mdctx.
    if mdctx initialization fails then
        Throw error and terminate.
    end if
    Begin SHAKE256 hashing process.
    Update hash with LHS, result.
    Finalize hash to produce hash_output$[N]$].
    for $i = 0$ to $N - 1$ do
        Extract $val$ from hash_output$[i]$.
        Compute $v[i] \leftarrow \mathrm{abs}(val) \mod Q1$.
        if $v[i] = 0$ then
            Set $v[i] \leftarrow 1$ to ensure non-zero component.
        end if
    end for
    Free SHAKE256 context: mdctx.
    Store $LHS$ in $C1$[lat].
end for
return $C1$.
```

Table 2: Parameter Values for Levels 1, 3, and 5

Level	**Dimension** (N)	ω ($R1$)	**Chain** (K)	($Q1$)	($SEED\_SIZE$)
1	64	81	8	257	16
3	128	9	6	257	24
5	256	3	4	257	32

Key Generation Algorithm

Algorithm 4 KeyGen()

Generates public and private keys.

```
Require: Prime modulus $Q1$, root $R1$, dimension $N$, number of chains $K$, seed size SEED_SIZE.
Initialize secretKeys[$K$][$N$] uniformly random $x$ in the range $[1, 255]$.
Initialize matrix: MATRIX_A[$K$][$N$][$N$].
Initialize matrix: MATRIX_U[$K$][$N$][$N$].
Generate random seed: PK_SEED_A of length SEED_SIZE.
Generate random seed: SK_SEED_U of length SEED_SIZE.
Sample MATRIX_A using PK_SEED_A.
Sample MATRIX_U using SK_SEED_U.
Initialize current_pk ← secretKeys[0].
current_pk ← NTT(current_pk, $Q1$, $R1$).
for I = 0 to $K - 1$ do          ▷ Iterate through the chain of transformations.
    Initialize result[$N$] ← 0.
    if I > 0 then
        Update skey ← secretKeys[I] and compute skey ← NTT(skey, $Q1$, $R1$).
        Element-wise multiplication: current_pk ← skey ∘ current_pk  mod $Q1$.
    end if
    result ← MatrixVectorProduct(MATRIX_A[I], current_pk, $Q1$).
    current_pk ← result.
    result[$N$] ← 0.
    result ← MatrixVectorProduct(MATRIX_U[I], current_pk, $Q1$).
    Apply hiding function: current_pk ← $pxz2$(result).
end for
Ensure non-zero condition: nonzero_count(current_pk) ≥ $N$.
if Condition fails then
    Retry key generation.
end if
return current_pk, PK_SEED_A, SK_SEED_U, secretKeys[K]
```

8.4 Signature Generation

Algorithm 5 Sign$(m, secretKeys[K], PK\_SEED\_A, pk\_elem, SK\_SEED\_U)$

Generates a signature for a message.

```
Require: Prime modulus $Q1$, root $R1$, dimension $N$, chain count $K$, seed size $SEED\_SIZE$, message $m$, , secret keys $secretKeys[K][N]$, public seed $PK\_SEED\_A$, public key $pk\_elem$, secret seed $SK\_SEED\_U$.
Initialize sig$[N] \leftarrow 0$.
Set SIG_COMPLETED $\leftarrow 0$.
Generate random bytes: $\text{rand\_A}, \text{rand\_B}$ of size $SEED\_SIZE$.
Compute FS using shake256$(\text{rand\_A}, \text{PK\_SEED\_A}, \text{pk\_element})$.
Compute SIG_SEED_B using shake256$(\text{FS}, m, \text{rand\_B}, \text{pk\_element})$.
Sample matrix: MATRIX_B$[K][N][N]$ using SIG_SEED_B.
Sample matrix: MATRIX_U$[K][N][N]$ using SK_SEED_U.
Initialize C1$[K][N]$ via genC1$(pk, \text{MATRIX\_B})$.
Set sig $\leftarrow$ secretKeys[0] and apply forward_ntt$(\text{sig}, Q1, R1)$.
for I = 0 to $K - 1$ do
    if I > 0 then
        skey $\leftarrow$ forward_ntt$(\text{secretKeys}[I], Q1, R1)$.
        sig $\leftarrow$ skey $\circ$ sig $\mod Q1$.
    end if
    result$[N] \leftarrow$ MatrixVectorProduct$(\text{MATRIX\_B}[I], \text{sig}, Q1)$.
    sig $\leftarrow$ result
    result$[N] \leftarrow 0$.
    result$[N] \leftarrow$ MatrixVectorProduct$(\text{MATRIX\_U}[I], \text{sig}, Q1)$.
    sig $\leftarrow$ sig $\circ$ C1[I] $\mod Q1$.
    Apply Hiding Function sig $\leftarrow pxz2(\text{result}))$.
end for
Apply inverse_ntt$(\text{sig}, Q1, R1)$.
Validate Entropy C3_CHECK $\leftarrow$ Verify_entropy$(sig)$.
C3 Check Retry after clearing buffers if C3_CHECK = 1.
Count non-zero elements in sig.
if nonzero_count(sig) $\geq$ N then
    Output $\text{sig}, \text{FS}, \text{rand\_B}$.
else
    Retry after clearing buffers.
end if
```

8.5 Signature Verification

This function uses the Fiat-Shamir heuristic to reconstruct the signature basis seed that was used during the signing process. Together with the public key matrix seed, both public and signature matrices are sampled and 'swapped', such that sig signature element is transformed by the public basis and the public key element is transformed by the signature basis. The public key is also isomorphically transformed by the masking value that was used at each layer during signing. This mask is derived from the interaction between the public key and the signature basis, effectively binding them together.

Additionally as valid signatures are information theoretically guaranteed to have observational entropy at or near maximum, we leverage this to detect potential forgeries.

Algorithm 6 Verify$(m, PK\_SEED\_A, pk\_elem, sig, FS, rand\_B)$

Verifies the signature of a message.

```
Require: Message $m$, public seed $PK\_SEED\_A$, public key $pk\_elem$, signature $sig$,
    Fiat-Shamir heuristic $FS$ and randomizer $rand\_B$.
Initialize SIG_SEED_B[$SEED\_SIZE$], C2[$N$], $C3\_CHECK \leftarrow 0$.
Validate Entropy C3_CHECK ← Verify_entropy($sig$).
C3 Check return 0 if C3_CHECK = 1.
Compute C2 using genC2($sig, pk\_elem, m, FS, rand\_B$).
Apply forward_ntt($C2, Q1, R1$) and forward_ntt($sig, Q1, R1$).
Construct temp_fs by concatenating FS, $m$, rand_B, and $pk\_elem$.
Compute SIG_SEED_B using shake256(temp_fs).
Sample matrix: PK_MATRIX[$K$][$N$][$N$] using PK_SEED_A
Sample matrix: SIG_MATRIX[$K$][$N$][$N$] using SIG_SEED_B.
Initialize LHS[$N$] ← $pk\_elem$, RHS[$N$] ← $sig$.
Compute C1[$K$][$N$] using genC1($pk, SIG\_MATRIX$).
for I = 0 to $K - 1$ do
    LHS ← LHS ∘ C2  mod $Q1$.
    LHS ← MatrixVectorProduct(SIG_MATRIX[I], LHS).
    LHS ← LHS ∘ C1[I]  mod $Q1$.
    if $I \neq K - 1$ then
        Apply Equivocation function: LHS ← $pxz2$(LHS).
    end if
end for
for I = 0 to $K - 1$ do
    RHS ← RHS ∘ C2  mod $Q1$.
    RHS ← MatrixVectorProduct(PK_MATRIX[I], RHS).
    if $I \neq K - 1$ then
        Apply Equivocation function: RHS ← $pxz2$(RHS).
    end if
end for
Compare RHS and LHS.
return 0 if equal, otherwise 1.
```

Observed Entropy Rejection Sampling

As part of what we are calling the third constraint, we want to rejection sample based on the observed randomness of the σ input. We are considering elements of length of $n = \{64, 128, 256\}$ and have implemented constraints thus far on two probabilistic features, bit level and byte level randomness. In the byte probability case, we are measuring each component byte value, with the ideal σ having zero colliding component values. However, we find that given the small sets of we are considering, we see multiple values appear two or three times, despite being valid.

To increase the granularity of randomness checking, we measure the raw ratio of 0 and 1 value bits across the array. For $n = 128$, where an ideally random signature would have

512 value 0 bits and 512 value 1 bits. To constrain input signatures to approximately 1/10 of the total possible modular space, we set the threshold ratio to 0.991. But, a "downshifted" element where components were in the range of $\{0, \ldots, 255\}$ entirely consisting of value 128 would pass with 512 0 bits and 512 1 bits. To correctly reject invalid signatures, we measure both bit level entropy and byte level entropy.

Table 3: Observed Entropy Thresholds for Different Values of using Byte probabilities N

N	**H_THRESHOLD**
64	5.8
128	6.8
256	7.8

Table 4: Observed Entropy Thresholds for Different Values of using Bit probabilities N

N	**HB_THRESHOLD**
64	0.991
128	0.991
256	0.991

Algorithm 7 Verify_Entropy(*sig*)

Validates the entropy of a signature.

```
Require: Signature sig.
Extract entropy-relevant components from sig: SIG_VALUES ← extract(sig).
Compute the empirical distribution DIST of SIG_VALUES over the modular domain [0, Q).
Calculate the Shannon entropy H_SIG:
    $$\text{H\_SIG} \leftarrow -\sum_{x \in \text{DIST}} p(x) \log p(x),$$
    where p(x) is the probability of x in DIST.
Extract count of 0 bits and 1 bits from sig as COUNT0 and COUNT1.
Compute Ratio as HB_SIG.
if H_SIG < H_THRESHOLD and HB_SIG < HB_THRESHOLD then
    return 1                                    ▷ Entropy too low, validation fails.
else
    return 0                                    ▷ Entropy validation succeeds.
end if
```

While the underlying concept of rejection sampling based on entropy should be clear, the exact implementation is to be refined in subsequent revisions of this preprint. The achievable constraint is that the adversary cannot forge σ using the entire ambient modular space, but only a specific fraction of it. This will lead to more refined and accurate probability.

9 Attack Models

In this section, we rigorously analyze the security of the proposed cryptographic scheme against both classical and quantum adversaries. We focus on proving that the scheme achieves IND-CPA (Indistinguishability under Chosen Plaintext Attack) security by demonstrating the computational infeasibility of recovering the hidden subgroup N or distinguishing ciphertexts under the specified attack models.

9.1 Classical Adversaries

Classical adversaries are limited to polynomial-time algorithms and lack quantum computational capabilities. We will show that, under standard cryptographic assumptions, such adversaries cannot feasibly recover the private keys or forge valid signatures.

9.1.1 Preliminaries

Let us recall the key components:

- The public key $\text{pk} = t'' = pxz2(t')$, where $t' = U \cdot t \mod q$ and $t = A \cdot x \mod q$.
- The mapping function $pxz2(\cdot)$ is a lossy, many-to-one function inducing high ambiguity.
- The hidden subgroup N is embedded in the non-abelian group $G = H_{\text{left}} \ltimes H_{\text{right}}$.

9.1.2 Proof of Security Against Classical Adversaries

Lemma 1 (Computational Indistinguishability). *Under the assumption that $pxz2(\cdot)$ is a pseudorandom function and that the underlying group operations are secure, any polynomial-time classical adversary has a negligible advantage in distinguishing between valid signatures and random elements, or in recovering the private key x or the matrix U.*

Proof. To prove this lemma, we proceed by contradiction. Assume there exists a polynomial-time classical adversary $\mathcal{A}$ that can distinguish valid signatures or recover x or U with non-negligible probability.

Step 1: Reduction to the Hardness of NAHSP.

Recall that recovering x or U is equivalent to solving the Non-Abelian Hidden Subgroup Problem (NAHSP) in the group G.

- The adversary's task reduces to finding N given oracle access to $f(g) = pxz2(U \cdot A \cdot x)$.

- As established in Section 3, solving NAHSP in this group is computationally infeasible for classical adversaries.

Step 2: Indistinguishability of $pxz2(\cdot)$ Outputs.

- The function $pxz2(\cdot)$ introduces high ambiguity, mapping exponentially many inputs to the same output.

- From Lemma 3 in Section 4.3, we know that the outputs of $pxz2(\cdot)$ are computationally indistinguishable from uniform random elements in $\mathbb{Z}_q^n$.

Step 3: Adversary's Advantage is Negligible.

- The adversary $\mathcal{A}$ cannot distinguish between $pxz2(U \cdot A \cdot x)$ and a random element without solving NAHSP.

- The probability that $\mathcal{A}$ successfully recovers x or U is bounded by $\varepsilon = \frac{1}{2^\lambda}$, where λ is the security parameter (e.g., $\lambda = 297$ as per the preimage count).
- Since ε is negligible, $\mathcal{A}$ cannot succeed with non-negligible probability.

Conclusion. Therefore, under standard cryptographic assumptions, no polynomial-time classical adversary can break the scheme, ensuring IND-CPA security against such adversaries.

9.2 Quantum Adversaries

Quantum adversaries have access to quantum computational resources, including algorithms like the Quantum Fourier Transform (QFT) and Grover's algorithm. We will demonstrate that even with these capabilities, adversaries cannot feasibly compromise the scheme.

9.2.1 Proof of Security Against Quantum Adversaries

Lemma 2 (Resistance to Quantum Attacks). *Under the assumption that the NAHSP is hard for quantum computers in non-abelian groups, and given the properties of the mapping function $pxz2(\cdot)$, any polynomial-time quantum adversary has a negligible advantage in breaking the scheme.*

Proof. Step 1: Non-Abelian Structure Prevents Efficient QFT-Based Attacks.
- Quantum algorithms like Shor's algorithm rely on the ability to perform efficient QFT over abelian groups.
- The group $G = H_{\text{left}} \ltimes H_{\text{right}}$ is non-abelian, as shown in Section 3.1.
- As a result, the standard QFT does not provide a means to solve the hidden subgroup problem efficiently in G.

Step 2: Ambiguity Introduced by $pxz2(\cdot)$.
- The mapping function $pxz2(\cdot)$ further complicates any attempt to extract information via quantum algorithms.
- From Lemma 7 in Section 4.7, even if a quantum adversary could invert $pxz2(\cdot)$, they would face an exponentially large preimage space, with 2^{297} indistinguishable candidates.

Step 3: Grover's Algorithm is Ineffective Due to Exponential Search Space.
- Grover's algorithm provides a quadratic speedup for unstructured search problems.
- Grover's algorithm generally is easily applied to chained systems with multiple secrets.

Step 4: No Known Quantum Algorithm Solves NAHSP in Non-Abelian Groups Efficiently.
- Despite extensive research, no quantum algorithm has been found that solves the NAHSP efficiently in general non-abelian groups.
- The hardness of NAHSP in such groups is a widely accepted assumption in quantum cryptography.

Conclusion. Given the non-abelian structure of the group and the properties of $pxz2(\cdot)$, quantum adversaries cannot break the scheme with non-negligible probability. Therefore, the scheme achieves IND-CPA security even in the presence of quantum adversaries.

9.3 IND-CPA Security Proof

We now provide a formal proof that the scheme is IND-CPA secure.

Theorem 1 (IND-CPA Security). *Under the assumption that the NAHSP is hard for both classical and quantum adversaries, and that $pxz2(\cdot)$ behaves as a pseudorandom function, the proposed digital signature scheme is IND-CPA secure.*

Proof. Definition of IND-CPA Security.

A digital signature scheme is IND-CPA secure if no polynomial-time adversary can distinguish between signatures of chosen messages, even when given access to a signing oracle.

Game-Based Proof Structure.

We consider the standard IND-CPA security game between a challenger and an adversary $\mathcal{A}$:

1. **Setup:** The challenger generates a public-private key pair (pk, sk) and provides pk to $\mathcal{A}$.

2. **Query Phase:** $\mathcal{A}$ may request signatures on messages of its choice.

3. **Challenge Phase:** $\mathcal{A}$ selects two messages m_0 and m_1. The challenger randomly selects $b \in \{0,1\}$ and returns $\sigma = \text{Sign}(m_b, sk)$.

4. **Guess Phase:** $\mathcal{A}$ outputs a guess b'. The adversary wins if $b' = b$.

Our goal is to show that $\Pr[b' = b] \leq \frac{1}{2} + \varepsilon$, where ε is negligible.

Analysis.

Assume, for contradiction, that $\mathcal{A}$ can win the game with a non-negligible advantage δ.

Step 1: Construction of a Simulator to Solve NAHSP.

We construct a simulator $\mathcal{S}$ that uses $\mathcal{A}$ to solve the NAHSP:

- $\mathcal{S}$ receives an instance of NAHSP in G and needs to find the hidden subgroup N.

- $\mathcal{S}$ simulates the challenger for $\mathcal{A}$, using the NAHSP instance to generate public keys and signatures.

- When $\mathcal{A}$ outputs b', $\mathcal{S}$ uses this information to extract information about N.

Step 2: Contradiction with the Hardness of NAHSP.

- If $\mathcal{S}$ can solve NAHSP using $\mathcal{A}$'s advantage δ, then the hardness assumption of NAHSP is violated.

- Therefore, δ must be negligible.

Step 3: Security Reduction via Hybrid Arguments.

- We can define a sequence of hybrid experiments transitioning from the real scheme to an ideal scheme where signatures are replaced with random values.

- The indistinguishability of outputs from $pxz2(\cdot)$ ensures that $\mathcal{A}$ cannot distinguish between hybrids with non-negligible advantage.

Conclusion.

Since any non-negligible advantage δ leads to a contradiction with the hardness of NAHSP, we conclude that $\mathcal{A}$ cannot win the IND-CPA game with more than negligible advantage. Therefore, the scheme is IND-CPA secure.

9.4 Resistance to Forgery Under Chosen Message Attacks

Theorem 2 (Unforgeability under Chosen Message Attack). *Assuming the hardness of NAHSP and the collision resistance of the hash function $J(\cdot)$, the proposed scheme is existentially unforgeable under chosen message attacks (EUF-CMA).*

Proof. **Step 1: Assumptions and Definitions.** - Let $\mathcal{A}$ be an adversary attempting to forge a valid signature σ^* for a message m^* that has not been queried during the signing oracle phase. - The scheme uses the Fiat-Shamir heuristic to bind the signature to the message, the public key, and a random nonce. - A valid forgery requires $\mathcal{A}$ to produce σ^* such that:

$$\text{Verify}(m^*, \sigma^*, pk) = \text{true}.$$

Step 2: Connection to NAHSP and $pxz2(\cdot)$. - To forge σ^*, $\mathcal{A}$ must either: 1. Recover the private key x or the hidden matrix U, allowing the computation of valid transformations. This is equivalent to solving the Non-Abelian Hidden Subgroup Problem (NAHSP), which is assumed to be hard. 2. Generate a valid preimage under $pxz2(\cdot)$ without access to the private key or matrix. The lossy, many-to-one nature of $pxz2(\cdot)$ ensures that the adversary cannot distinguish valid preimages from an exponentially large indistinguishable set.

Step 3: Resistance to Hash Function Collisions. - The Fiat-Shamir heuristic involves the hash function $J(\cdot)$, which produces a binding challenge for the signature. For $\mathcal{A}$ to forge σ^*, it must either: 1. Find a collision $J(pk \parallel m^* \parallel r) = J(pk \parallel m' \parallel r')$, which is infeasible due to the assumed collision resistance of $J(\cdot)$. 2. Guess the challenge generated by $J(\cdot)$ and align it with a valid subgroup element. The probability of such a guess is negligible due to the high entropy of the output space of $J(\cdot)$.

Step 4: Reduction to a Hard Problem. - Assume $\mathcal{A}$ successfully forges σ^* with non-negligible probability. We construct a simulator $\mathcal{S}$ that uses $\mathcal{A}$ to solve NAHSP or find a collision in $J(\cdot)$: 1. $\mathcal{S}$ simulates the signing oracle for $\mathcal{A}$, generating signatures using a secret key x and the private matrix U. 2. If $\mathcal{A}$ outputs a valid forgery σ^*, $\mathcal{S}$ uses σ^* to extract information about the hidden subgroup N or to find a collision in $J(\cdot)$. 3. Since both outcomes contradict the hardness of NAHSP or the collision resistance of $J(\cdot)$, $\mathcal{A}$'s success probability must be negligible.

Step 5: Conclusion. - The adversary $\mathcal{A}$ cannot forge a valid signature σ^* on a message m^* without solving NAHSP, inverting $pxz2(\cdot)$, or finding a collision in $J(\cdot)$, all of which are computationally infeasible. - Therefore, the proposed scheme is existentially unforgeable under chosen message attacks (EUF-CMA).

9.5 Inapplicability of CCA2 Security

It is important to note that oue proposed digital signature scheme does not incorporate a decryption oracle, as it is not designed to handle encrypted messages or ciphertext directly. The absence of such an oracle renders the chosen ciphertext attack (CCA2) model irrelevant for this construction.

Instead, the security of the scheme is analyzed under the IND-CPA (Indistinguishability under Chosen Plaintext Attack) and EUF-CMA (Existential Unforgeability under Chosen Message Attack) models, which are sufficient and appropriate given the nature of the signature application.

By excluding a decryption oracle, the scheme eliminates a common attack vector associated with adaptive adversaries in CCA2 scenarios, further solidifying its robustness in practical cryptographic deployments.

9.5.1 Brute Force Key Recovery

Brute force recovery of secrets is implementation dependent. The scheme covered in this document leverages a single private matrix seed for each instance in the chain, and unique password x elements per instance. Based on size and security concerns, if private key size was critical, the private key, in theory, could be shrunk to a single secret seed that expanded to provide every hidden matrix and x element. If absolute security were paramount, each hidden matrix could be derived from its own seed, or stored fully instantiated. In this brief analysis, we will simply derive the costs of brute forcing the scheme as described. Each level has k chains, with one x secret element, effectively n bytes long. Additionally, each level has one hidden matrix seed, of $SEED\_SIZE$ length, in bytes. Relative sizes are listed below:

Table 5: Brute Force Secret Byte Analysis

n	k	Single x	All x	Hidden Seed	Total Secret	Complexity
64	8	64	512	16	528	2^{4224}
128	6	128	768	24	792	2^{6336}
256	4	256	1024	32	1056	2^{8448}

As even level I requires 2^{4224} classical operations, brute force attacks do not appear to be a practical concern with the variant as described in this paper. Note that this could change depending on various implementation optimizations.

9.6 Conclusion on Attack Models

Through rigorous proofs, we have established that the proposed cryptographic scheme is secure against both classical and quantum adversaries. The security relies on:

- The computational hardness of the NAHSP in non-abelian groups.
- The pseudorandomness and computational indistinguishability introduced by the mapping function $pxz2(\cdot)$.
- The collision resistance of the hash function used in the Fiat-Shamir heuristic.

These properties collectively ensure that adversaries cannot feasibly recover private keys, forge signatures, or distinguish ciphertexts, thereby achieving IND-CPA security and resisting forgery under chosen message attacks.

10 Implementation and Efficiency

10.1 Performance Evaluation

Table 6: Compute Cycles (in Megacycles) for Key Generation, Signing, and Verification

Level	Key Generation (Mc)	Signature (Mc)	Verification (Mc)
I	.49	.38	.44
III	1.03	.958	1.10
V	9.46	3.25	3.48

Platform: Apple MacBook M2 MAX with 32 GB RAM.

Level	n	PK (bytes)	Sig (bytes)	k
I	64	80	96	8
III	128	152	176	6
V	256	288	320	4

Table 7: Public Key, Signature Sizes, and Chain Instances Across Levels

10.2 Comparison with Other Schemes

Table 8: Performance Metrics of Alternative Signature Algorithms (in Bytes and Megacycles)

Algorithm	PK+SIG (Bytes)	Sign (Mcycles)	Verify (Mcycles)
Dilithium2	3,732	0.333	0.118
Dilithium3	5,261	0.529	0.179
Dilithium5	7,219	0.642	0.280
MAYO1	1,489	.461	.175
MAYO3	3,233	1.664	.610
MAYO5	5,846	4.150	1.186
HAWK-512	1,573	.085	.148
HAWK-1024	3,661	.180	.303
Falcon-512	1,563	1.010	.081
Falcon-1024	3,073	2.053	.161
SLH-DSA-128f	17,120	239.794	12.910
SLH-DSA-192f	35,712	386.862	19.877
SLH-DSA-256f	49,920	763.942	19.886
SQIsign-1	241	5,669.00	108.00
SQIsign-3	359	43,760.00	654.00
SQIsign-5	463	158,544.00	2,177.00
EQISIGN-I	176	.38	.44
EQISIGN-III	328	.958	1.10
EQISIGN-V	608	3.25	3.45

The above data was been gathered from the PQShield Post-Quantum signatures zoo, we have not verified it and the most recent developments may not be reflected, but we feel the source is accurate for this early comparison. To date, our priority has been theoretical security and communication efficiency, with little attention paid to performance optimization. Over time, it is almost certain our performance numbers will improve beyond our reference instance. We are using portable ANSI C, openSSL/TSL, our portion of the code does not leverage intrinsics, is single threaded, and the majority of time is currently spent expanding matrices from seed. With dedicated effort we feel performance and optimization will yield significant improvements. That said, we do not expect to outperform ML-DSA in terms of compute, nor do we expect computational performance to be a barrier for adoption.

10.3 Rejection Sampling of Zero Coefficients in Output Variables

The elimination of zero value elements is a strong method to create resilience against heuristic cryptanalysis, both quantum and classical. Due to the zero product property of finite fields, allowing the value zero tends to cause accumulation in output variables (keys, signatures) opening a window for exploitation. Additionally, while not covered elsewhere, in practical implementations, combined with the working modulus of $q = 257$, we are able to leverage the possible component range to increase communications efficiency. Under normal circumstances, operating modulo 257 results in elements that range from $\{0, ..., 256\}$, or 257 unique values, requiring 9 bits to accurately represent. The elimination of zero via rejection sampling allows each value to map to 256 elements which can be represented using 8 bits.

Simply, when serializing variables for transmission, before transmission we simply subtract one from each value. Upon receiving keys or signatures we 'reconstitute' them by incrementing each by one, mapping back to the appropriate original values. This is an easy optimization to align with normal machine word boundaries.

11 Open Research Questions

1. For function pxz2(), the optimal ratio of class size to class number in relation to group sizes $|G|$ and $|N|$ is an open question.

2. Correct formal complexity classification of NAHSP under extreme equivocation. The NAHSP is conjectured to be in the EXP complexity class, with no known efficient quantum algorithm. These conjectures should be proven if possible, but this work is outside the scope of this paper.

12 Applications of Matrix based NAHSP-Based Cryptography

The NAHSP-based cryptographic scheme offers a versatile, lightweight, and quantum-resistant framework for diverse applications. Its compact signatures, efficient communica-

[0] Data source: `https://pqshield.github.io/nist-sigs-zoo/c`

tion requirements, and ability to deploy through software patches without new hardware set it apart from traditional lattice-based approaches such as Dilithium. These features enable its use in domains ranging from terrestrial networks to undersea and RF-limited environments, making it uniquely suited for next-generation cryptographic needs.

12.1 Core Cryptographic Capabilities

An NAHSP-based scheme based on bilinear matrices can be extended beyond foundational primitives to more advanced cryptographic constructions:

- Digital Signatures: Compact and efficient signatures ensure secure authentication, document signing, and certificate management, with communication sizes significantly smaller than lattice-based systems like Dilithium.
- Public Key Agreement: Enables fast, quantum-resilient key exchanges with minimal communication overhead, ideal for constrained networks.
- Identity-Based and Attribute-Based Cryptography: Supports fine-grained access control, allowing secure communication based on user identities or attributes without requiring heavy key distribution infrastructure.
- Zero-Knowledge Proofs (ZKPs): Facilitates privacy-preserving verification of statements without exposing underlying secrets, essential for regulatory compliance and secure interactions.

12.2 Efficient Communication and Deployability

Compact Communication Sizes:

- NAHSP-based cryptography achieves extremely small communication footprints, with signatures and keys often requiring a fraction of the size used by lattice-based systems. This efficiency is critical in bandwidth-limited environments such as undersea and RF networks.
- Example: For a security level equivalent to Dilithium-III, the NAHSP-based scheme offers signatures of 80 bytes and public keys under 100 bytes, compared to the hundreds or thousands of bytes required by Dilithium.

Software-Only Deployment:

- Unlike lattice-based systems, which often require specialized hardware for efficient operation, NAHSP-based cryptography can be implemented as a drop-in replacement via software patches.
- This allows immediate deployment in existing infrastructures, including mobile devices, routers, and IoT systems, without the need for hardware upgrades.
- Rapid updates ensure forward compatibility with evolving security standards while minimizing deployment costs.

12.3 Secure Communication Across Diverse Environments

NAHSP-based cryptography's compact and efficient design enables secure communication in challenging and bandwidth-constrained domains:

- Cellular Networks: Ensures efficient and secure handshakes, even in low-latency 5G environments, where minimal communication overhead is critical.
- Radio Frequency (RF) Systems: Compact key sizes and signatures reduce transmission time in RF-constrained settings, such as military radios and satellite uplinks.
- Undersea Acoustics: Low-bandwidth undersea acoustic networks benefit from NAHSP's compact communication, enabling secure exchanges where data rates are severely limited.

12.4 Comparison with Dilithium and Other Systems

A matrix based NAHSP scheme addresses several limitations of Dilithium and similar lattice-based approaches:

- Smaller Communication Sizes: Signatures and keys are significantly more compact, reducing storage and bandwidth requirements.
- Flexibility Across Environments: Performs robustly in environments where lattice-based schemes face challenges, such as RF and undersea communication.
- Advanced Constructions: Offers natural support for identity-based encryption and zero-knowledge proofs, features that are generally not feasible to implement over schemes with noise.

12.5 Real-World Applications

Critical Infrastructure and Defense:

- Secures command and control systems in military and intelligence operations, ensuring quantum resilience and adaptability across RF and satellite links.
- Protects SCADA systems in critical infrastructure, such as energy grids and transportation networks, with lightweight and efficient cryptographic primitives.

IoT and Edge Devices:

- Provides secure authentication for IoT devices with constrained processing power, mitigating risks of botnet attacks and data breaches.
- Ensures efficient encryption and signing for edge devices in industrial and healthcare settings.

Blockchain and Distributed Systems:

- Enhances consensus mechanisms with compact, quantum-resistant signatures, reducing energy consumption and improving scalability.

- Secures smart contracts and cryptographic tokens with lightweight, efficient constructions.

Telecommunications and Financial Systems:

- Enables secure mobile payment systems and digital banking with minimal communication overhead, ensuring transaction authenticity and integrity.
- Modernizes public key infrastructure (PKI) for quantum resilience while minimizing deployment costs.

12.6 Conclusion

While the primary embodiment of this invention is a digital signature scheme leveraging the Non-Abelian Hidden Subgroup Problem (NAHSP) and equivocation via the pxz2() mapping function, however the core mechanism is broadly applicable to other cryptographic primitives. These include, but are not limited to, public key exchange, encryption schemes (such as identity-based and attribute-based encryption), and zero-knowledge proofs. The underlying NAHSP-based obfuscation provides a foundation for secure and efficient cryptographic systems across various applications.

13 Future Work and Concluding Remarks

- While theoretically robust, the construction requires careful selection of optimal variables and implementation of the underlying mathematics to run in constant time. Achieving constant time execution mitigates side-channel attack and aligns with best practices for any algebraic cryptographic scheme.
- Like all novel forms of cryptography, extensive adversarial cryptanalysis is required. We welcome experienced cryptanalysis focused collaboration.
- Optimization of functions to leverage platform-specific SIMD instructions can significantly accelerate operations while maintaining constant runtime guarantees. This will enhance the scheme's practical viability across diverse hardware platforms.
- The scheme employs rejection sampling on intermediates with zero-value coefficients to prevent degeneracy. A thorough adversarial analysis is needed to evaluate whether this rejection sampling introduces potential vulnerabilities that could aid cryptanalysis. If identified, appropriate mitigations must be developed.
- Additionally, the rejection sampling method of zero components is fairly simple and is currently a major performance cost. By implementing a more optimal mechanism, these costs can be minimized.

Integrating quantum-resilient cryptographic systems into existing infrastructures remains a complex challenge, particularly for hardware-constrained environments or legacy systems reliant on established PKI frameworks. This NAHSP-based system offers a promising approach by providing compact signatures and practical efficiency suitable for retrofitting into current infrastructures, including standard-sized X.509 certificates.

These properties make it a strong candidate for addressing the scalability and trust requirements needed in the transition to post-quantum security.

While this work may be among the first asymmetric cryptographic systems to aim for practical information-theoretic security guarantees by design, its formal proofs and construction serve as a foundational step toward bridging the gap between theoretical resilience and real-world application. This approach diversifies the cryptographic landscape, complementing existing quantum-resilient efforts and enhancing robustness against diverse attack vectors.

Beyond its immediate applications, this cryptosystem opens new avenues of research across multiple fields:

- Cryptography: The NAHSP framework invites exploration into additional constructions such as group-based encryption, secure multi-party computation, and advanced privacy-preserving protocols.
- Complexity Theory: By leveraging non-abelian group properties, the system provides fertile ground for studying alternate hardness assumptions and their implications for classical and quantum computational limits.
- Quantum Algorithm Design: The inherent resilience of the scheme challenges researchers to explore novel quantum algorithms capable of addressing non-abelian group problems, advancing our understanding of quantum computational power.
- Systems Security: With its adaptability to constrained environments such as IoT, RF, and undersea acoustics, this system sets the stage for breakthroughs in secure communication under extreme conditions.

Closing Statement: This cryptographic scheme offers a significant contribution to the evolving landscape of post-quantum security. By providing practical information-theoretic guarantees and addressing key implementation challenges, it has the potential to transform how secure systems are designed and deployed. While further research and optimization remain, this work lays a strong foundation for future innovations in cryptography, complexity theory, and quantum algorithm design, positioning it as a critical component in the journey toward resilient and scalable global security systems.

14 Acknowledgements

This work is provisionally patented, USPTO 63/726,609.

References

[1] Daniel J. Bernstein et al. “The SPHINCS+ Signature Framework”. In: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’19. London, United Kingdom: Association for Computing Machinery, 2019, pp. 2129–2146. ISBN: 9781450367479. DOI: 10.1145/3319535.3363229. URL: https://doi.org/10.1145/3319535.3363229.

[2] Léo Ducas et al. “Dilithium: A high-speed lattice-based digital signature scheme”. In: *CCS 2019*. 2019, pp. 897–918.

[3] Luca De Feo et al. *SQISign: compact post-quantum signatures from quaternions and isogenies.* Cryptology ePrint Archive, Paper 2020/1240. 2020. URL: `https://eprint.iacr.org/2020/1240`.

[4] David Garber. "Braid group cryptography". In: *Braids: Introductory lectures on braids, configurations and their applications.* World Scientific, 2010, pp. 329–403.

[5] Dimitri Grigoriev and Ilia Ponomarenko. *Constructions in public-key cryptography over matrix groups.* 2005. arXiv: `math/0506180 [math.GR]`. URL: `https://arxiv.org/abs/math/0506180`.

[6] Ki Hyoung Ko et al. "New public-key cryptosystem using braid groups". In: *Advances in Cryptology—CRYPTO 2000: 20th Annual International Cryptology Conference Santa Barbara, California, USA, August 20–24, 2000 Proceedings 20.* Springer. 2000, pp. 166–183.

[7] Patrick Longa, Wen Wang, and Jakub Szefer. *The Cost to Break SIKE: A Comparative Hardware-Based Analysis with AES and SHA-3.* Cryptology ePrint Archive, Paper 2020/1457. 2020. URL: `https://eprint.iacr.org/2020/1457`.

[8] Karl Mahlburg. "An overview of braid group cryptography". In: *preprint* (2004).

[9] Alexei Myasnikov, Vladimir Shpilrain, and Alexander Ushakov. "A practical attack on a braid group based cryptographic protocol". In: *Advances in Cryptology–CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005. Proceedings 25.* Springer. 2005, pp. 86–96.

[10] Alexei G Myasnikov, Vladimir Shpilrain, and Alexander Ushakov. *Non-commutative cryptography and complexity of group-theoretic problems.* 177. American Mathematical Soc., 2011.

[11] Michael Schmid et al. *Falcon Takes Off - A Hardware Implementation of the Falcon Signature Scheme.* Cryptology ePrint Archive, Paper 2023/1885. 2023. URL: `https://eprint.iacr.org/2023/1885`.

[12] Claude E. Shannon. "Communication theory of secrecy systems". In: *Bell Syst. Tech. J.* 28.4 (1949), pp. 656–715. DOI: `10.1002/J.1538-7305.1949.TB00928.X`. URL: `https://doi.org/10.1002/j.1538-7305.1949.tb00928.x`.

[13] Claude Elwood Shannon. "A Mathematical Theory of Communication". In: *The Bell System Technical Journal* 27 (1948), pp. 379–423. URL: `http://plan9.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf` (visited on 04/22/2003).