

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

4 Samuel Lavery

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

8 December 27, 2024

9 **Abstract**

10 We present a novel digital signature scheme grounded in non-commutative cryp-
11 tography and implemented over a bilinear matrix group platform. At the core
12 of our design is a unique equivocation function that obfuscates intermediate ele-
13 ments, effectively concealing outputs and minimizing observable information leak-
14 age. To the best of our knowledge, this is the first digital signature scheme to
15 combine information-theoretic security with computational hardness, relying on a
16 challenging instance of the Non-Abelian Hidden Subgroup Problem (NAHSP) and
17 strengthened by practical guarantees. This dual-layered security approach ensures
18 robustness against both classical and quantum adversaries while maintaining com-
19 munication overheads competitive with RSA. Our work represents a significant
20 advancement toward efficient, quantum-resilient digital signatures for real-world
21 applications. This paper is an early pre-release intended to invite collaboration
22 and feedback. The work is presented for research purposes only and is not intended
23 for use in production systems.

24 **Contents**

25 **1 Introduction** 4
26 1.1 The State of the Art 4
27 1.2 Secure and Efficient Signatures from the Non-Abelian Hidden Subgroup
28 Problem and Information Theory 4
29 **2 High-Level Description** 6
30 2.1 Core Representation 6
31 2.2 Construction 7
32 2.2.1 Definition of the Group G and Subgroups H_{right} and H_{left} 7
33 **3 Problem Statement** 10

34	4 Preliminary Results and Contributions	15
35	4.1 Structure of Remainder of Paper	18
36	5 Notation and Definitions	18
37	6 Formal Reduction to the Non-Abelian Hidden Subgroup Problem (NAHSP)	19
38	6.1 Group Structure and Properties	19
39	6.2 Formal Definition of NAHSP	19
40	6.3 Oracle Construction and Reduction to NAHSP	20
41	6.4 Hardness of Subgroup Recovery	20
42	6.5 Reduction	21
43	7 Equivocation and Indistinguishability from $px(\cdot)$	22
44	7.1 Equivocation Function $px(\cdot)$	22
45	7.2 Analysis of $px(\cdot)$	24
46	7.3 Probability of Random Forgery	24
47	7.3.1 Brute Force Storage Requirements	25
48	7.4 Equivalence Class Invariance of $px(\cdot)$	26
49	7.5 Enhanced Hiding of $px(\cdot)$	27
50	7.6 Indistinguishability and Statistical Uniformity of $px(\cdot)$ Outputs	28
51	7.7 Collision and Pre-Image Resistance of $px(\cdot)$	30
52	7.8 Avalanche Effect in $px(\cdot)$	31
53	7.9 Adaptive Security of $px(\cdot)$	32
54	7.10 Adversarial Complexity and Relation to NAHSP	33
55	7.11 Quantum Resistance of $px(\cdot)$	35
56	8 Information-Theoretic Security of $px(t')$	36
57	9 Proof of Consistency as Verification Under Homomorphic Transformations	37
58		
59	10 Implementation Details	39
60	10.1 Matrix Generation Using Diverse Cryptographically Secure PRNGs	39
61	10.1.1 Public Matrix Generation	39
62	10.1.2 Private Matrix Generation	39
63	10.1.3 Security Implications	40
64	10.2 Algorithm Details	40
65	10.3 Utility Algorithms	40
66	10.4 Signature Generation	44
67	10.5 Signature Verification	44
68	10.6 Hiding Function	47
69	11 Attack Models	48
70	11.1 Classical Adversaries	49
71	11.1.1 Preliminaries	49
72	11.1.2 Proof of Security Against Classical Adversaries	49
73	11.2 Quantum Adversaries	50
74	11.2.1 Proof of Security Against Quantum Adversaries	50
75	11.3 IND-CPA Security Proof	50

76	11.4 Resistance to Forgery Under Chosen Message Attacks	51
77	11.5 Inapplicability of CCA2 Security	52
78	11.5.1 Brute Force Key Recovery	52
79	11.6 Conclusion on Attack Models	53
80	12 Implementation and Efficiency	53
81	12.1 Performance Evaluation	53
82	12.2 Comparison with Other Schemes	54
83	12.3 Rejection Sampling of Zero Coefficients in Output Variables	55
84	13 Open Research Questions	55
85	14 Applications of Matrix based NAHSP-Based Cryptography	55
86	14.1 Core Cryptographic Capabilities	55
87	14.2 Efficient Communication and Deployability	56
88	14.3 Secure Communication Across Diverse Environments	56
89	14.4 Comparison with Dilithium and Other Systems	57
90	14.5 Real-World Applications	57
91	14.6 Conclusion	58
92	15 Future Work and Concluding Remarks	58
93	16 Acknowledgements	59

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], a bit of 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 is still 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 generally by algebraic cryptanalysis. With most schemes having core algebraic and structural weaknesses[9], the bulk of research inertia has 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

138 cryptography[5], the chosen matrices were structured, lacked entropy, had exploitable sub-
139 groups, or exploitable linear relationships. These schemes relied on matrix equivalence
140 or conjugacy problem hardness assumptions, which turned out to be less robust than
141 anticipated. Our work uses a pair of dense, full rank random matrices, forming a bilin-
142 ear group. By using non-correlated matrices, we mitigate attacks leveraging structural
143 relationships. With half of the matrix group randomized each operation, observed corre-
144 lation only applies to a single matrix pair. This acts to effectively minimize exploitable
145 relationship information. This bilinear matrix group structure, combined with a novel
146 equivocation function form the building blocks of our instance of the non-Abelian Hidden
147 Subgroup Problem (NAHSP).

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

157 In the context of this construction, the low dimensions $n = \{64, 128, 256\}$ and small
158 modulus size $q = 257$ do not directly result in an intractable NAHSP instance. Nor-
159 mally, an adversary observing a small number of intermediate outputs could construct
160 a solvable system of linear equations via modular arithmetic and Gaussian elimination.
161 However, randomization acts to uniformly project intermediate output elements of matrix
162 calculations across a large ambient module space \mathbb{Z}_n^q . We leverage information-theoretic
163 principles[13], equivocation and mutual entropy, to prove practical information theoretic
164 security guarantees. While we do not claim the perfect secrecy of the one-time pad[12],
165 we use information theoretic arguments to quantify the computational infeasibility of
166 reconstructing the problem. We employ a novel equivocation function to partition the
167 ambient space into equivalence classes, dispersing secret entropy across a class of indistin-
168 guishable valid pre-images. If the original output is the needle, the adversary is given the
169 haystack. This core disambiguation ensures that each equivalence class is populated with
170 indistinguishable pre-images, obfuscating the relationship between inputs and outputs.
171 This indistinguishability enables the transformation our computationally hard matrix
172 problem into an information theoretic one. Thus, the challenge for an adversary shifts
173 from using linear algebra to solve a well posed system of equations to identifying the
174 correct indistinguishable elements needed to construct it.

175 This work employs a hybrid security model that amplifies computational complexity
176 with information-theoretic security. By leveraging uniform distribution and indistin-
177 guishability, we transform the computational problem of solving the NAHSP into an
178 information-theoretic challenge. The resilience of the NAHSP to quantum attack at-
179 tempts underscores its suitability as a basis for this form of construction.

180 In practical terms, this work challenges the prevailing notion that achieving quantum
181 resilience necessitates a significant increase in communication overhead. Our construc-
182 tion integrates the structural advantages of non-abelian matrix groups with information-
183 theoretic principles, offering an alternative direction for designing cryptographic systems
184 that balance quantum resistance with real-world usability. The short signatures and pub-
185 lic keys comparable to ECDSA and RSA may be especially useful in highly constrained

186 network environments, where more established lattice based schemes struggle to integrate.
 187 This work is a preprint released to facilitate discussion and collaboration. Updates
 188 and refinements will follow as necessary.

189 2 High-Level Description

190 2.1 Core Representation

191 Our instance of the Non-Abelian Hidden Subgroup Problem (NAHSP) is defined within
 192 the context of a bilinear matrix group G . The group operation, matrix-vector multipli-
 193 cation, involves a public matrix A and a secret vector x , while the hidden subgroup N is
 194 implicitly defined through transformations induced by the secret matrix U . Specifically,
 195 the outputs $t \equiv Ax \pmod{q}$ and $t' \equiv Ut \pmod{q}$ are observed, with U and x remaining
 196 hidden. To amplify cryptographic hardness, the equivocation function px maps these
 197 outputs into indistinguishable equivalence classes, obfuscating the relationship between
 198 A , U , x , and the observed results.

199 Our variant can be concisely represented for key generation as:

$$\begin{aligned} & t \equiv A \cdot x \pmod{q}, \\ & t' \equiv U \cdot t \pmod{q}, \\ & pk = t'' \equiv px(t') \pmod{q}, \end{aligned}$$

202 For signing as:

$$\begin{aligned} & t \equiv B \cdot x \pmod{q}, \\ & t' \equiv U \cdot t \pmod{q}, \\ & sig = t'' \equiv px(t' \circ J(C1)) \pmod{q}, \end{aligned}$$

205 For validation as:

$$\begin{aligned} & LHS \equiv B \cdot (pk \circ J(C2)) \pmod{q}, \\ & LHS' \equiv px(LHS \circ J(C1)) \pmod{q}, \\ & RHS \equiv A \cdot (sig \circ J(C2)) \pmod{q}, \\ & RHS' \equiv px(RHS) \pmod{q}, \end{aligned}$$

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

207 where:

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

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

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

211 x : Secret vector uniformly sampled from $\mathbb{Z}_{1-256}^n = \{1, 2, \dots, 256\}^n$, ensuring no zero
 212 entries.

- 213 t : Element in the right subgroup $H_{\text{right}} \subseteq G$, spanning the ambient space and serving
 214 as a hidden input to H_{left} .
- 215 $J()$: Secure hash function (e.g., SHA3/SHAKE) producing pseudorandom outputs in
 216 \mathbb{Z}_q .
- 217 t' : Element in the hidden left subgroup $N \subseteq G$.
- 218 $px()$: Mapping function projecting t' into equivalence classes, ensuring computational
 219 infeasibility of N 's recovery.
- 220 t'' : Obfuscated output element in the superset $N' \subseteq G$.
- 221 q : Prime modulus defining the finite field \mathbb{Z}_q .
- 222 fs : Fiat-Shamir heuristic in \mathbb{Z}_q^n , binding the public key, message, and randomness to
 223 the signature context.
- 224 r : Signature randomizer value in \mathbb{Z}_q^n , mitigating replay attacks and enhancing security
 225 against sEU-CMA.
- 226 pk : Public key element in \mathbb{Z}_q^n .
- 227 $C1$: Forgery constraint 1, an element in \mathbb{Z}_q^n , derived as a hash of intermediate context
 228 values related to $pk \cdot B$.
- 229 $C2$: Forgery constraint 2, an element in \mathbb{Z}_q^n , derived as a hash of pk , sig , and other
 230 context elements.
- 231 \circ : Element-wise multiplication operation.
- 232 \cdot : Matrix-vector product operation.

233 2.2 Construction

234 2.2.1 Definition of the Group G and Subgroups H_{right} and H_{left}

235 Ambient Group G_{ambient} :

236 The ambient modular group defines the space of possible elements as:

$$G_{\text{ambient}} = \mathbb{Z}_q^n,$$

237 representing the set of all n -dimensional vectors over \mathbb{Z}_q . The group operation is matrix-
 238 vector multiplication modulo q . This expansive space provides a form of limited one-time
 239 pad security, where the inherent size and approximate uniform partitioning of G_{ambient}
 240 by $px(\cdot)$ ensure that individual elements are computationally indistinguishable without
 241 knowledge of the specific subgroup structure.

242 Working Group G :

243 The working group G in the context of the NAHSP is defined as the group generated by
 244 the subgroups H_{right} and H_{left} :

$$G = \langle H_{\text{right}}, H_{\text{left}} \rangle,$$

245 where each subgroup introduces hidden structure critical to cryptographic security.

246 **Subgroup H_{right} :**

247 The subgroup H_{right} generates the space of input elements t based on the public matrix
248 A :

$$H_{\text{right}} = \{t \mid t = A \cdot x \pmod{q}, x \in \mathbb{Z}_q^n\},$$

249 where:

- 250 • $A \in \mathbb{Z}_q^{n \times n}$: Dense, full-rank, random public matrix.
- 251 • $x \in \mathbb{Z}_q^n$: Secret vector sampled uniformly from $\{1, \dots, q-1\}^n$.

252 **Subgroup H_{left} :**

253 The subgroup H_{left} maps elements t from H_{right} through the private matrix U :

$$H_{\text{left}} = \{t' \mid t' = U \cdot t \pmod{q}, t \in H_{\text{right}}\},$$

254 where:

- 255 • $U \in \mathbb{Z}_q^{n \times n}$: Hidden, dense, full-rank private matrix.
- 256 • $t \in H_{\text{right}}$: Hidden input generated by the public matrix A .

257 **Hidden Subgroup N :**

258 The hidden subgroup N is a normal subgroup embedded within H_{left} :

$$N = \{t' \in H_{\text{left}} \mid t' \cdot t^{-1} \in H_{\text{left}}, \forall t \in H_{\text{right}}\}.$$

259 The structure of N ensures the following cryptographic properties:

- 260 • **Normality and Algebraic Consistency:** The subgroup N maintains its normality
261 within H_{left} , preserving the relationship $gN = Ng$ for all $g \in G$. This is essen-
262 tial for reductions to the Non-Abelian Hidden Subgroup Problem (NAHSP).
- 263 • **Obfuscation by Subgroup Relationships:** Elements of N are indistinguishable from
264 non-members without knowledge of the private transformations U and x , as the
265 coset structure of N is obscured by the combined transformations within H_{left} .
- 266 • **Cryptographic Hardness:** Recovering N requires solving the NAHSP, which is com-
267 putationally infeasible under both classical and quantum adversarial models.

268 **Role of the Hidden Matrix U :**

269 The hidden matrix U introduces non-commutative transformations that amplify crypto-
270 graphic hardness by ensuring that the output space spans the full ambient group:

- 271 • **Full-Rank Transformation:** The full-rank nature of U guarantees that the transfor-
272 mation $t' = U \cdot t \pmod{q}$ spans the entire output space \mathbb{Z}_q^n . This ensures that the
273 final output is not constrained by structural dependencies, maximizing the entropy
274 of t' .

275 • Non-Commutative Action: The transformation $U \cdot t \pmod q$ introduces non-commutative
276 operations, breaking linear relationships and obfuscating the structure of H_{left} . This
277 makes it computationally infeasible for an adversary to recover U or t without solv-
278 ing the underlying subgroup problem.

279 • Preservation of Subgroup Properties: The transformation maintains the algebraic
280 properties of H_{left} and preserves the coset structure of the hidden subgroup N
281 under conjugation:

$$gN = Ng, \quad \forall g \in G.$$

282 This preservation is critical for ensuring cryptographic consistency and enabling
283 reductions to the NAHSP.

284 **Role of the Public Matrix A and Secret Vector x :**

285 The public matrix A and the secret vector x jointly ensure that transformations are
286 randomized and uniformly distributed across the input space, enabling the outputs to
287 fully span the group \mathbb{Z}_q^n . Together, A and x mitigate adversarial attacks by introducing
288 fresh randomness and structural complexity for every operation:

289 • Full-Rank Input Transformation: The full-rank nature of A ensures that $t = A \cdot$
290 $x \pmod q$ spans the entire space of valid inputs, \mathbb{Z}_q^n . This guarantees that every
291 operation begins with a uniformly distributed element, eliminating structural biases.

292 • Per-Operation Randomization: The randomized public matrix A , combined with
293 the secret vector x , ensures that transformations differ across operations. This
294 randomness prevents correlation attacks and ensures that adversaries cannot infer
295 relationships between consecutive operations.

296 • Obfuscation of Intermediate Values: The secret vector x acts as a one-time secret
297 for each transformation, ensuring that the intermediate value t remains private and
298 uncorrelated with U .

299 • Resistance to Chosen Message Attacks: The per-operation uniqueness of A and x
300 ensures that information observed from one operation cannot be reused or leveraged
301 to attack subsequent operations. This protects the scheme from signature re-use
302 attacks.

303 **Security Mechanisms:**

304 The security of the scheme is rooted in the computational hardness of solving the Non-
305 Abelian Hidden Subgroup Problem (NAHSP):

306 • Computational Hardness: Recovering the hidden subgroup N or reconstructing U
307 and x from observed outputs requires solving the NAHSP, a problem resistant to
308 both classical and quantum adversaries.

309 • Structural Complexity: The layered transformations through A and U , combined
310 with the subgroup relationships between H_{right} , H_{left} , and N , amplify the inherent
311 difficulty of the problem.

312 **Adversary’s Computational Task:**

313 The adversary’s goal is to recover N and reconstruct U and x . This task is rendered
314 cryptographically non-trivial due to:

- 315 • Full-Rank Transformations: The full-rank nature of A and U ensures that observed
316 outputs span the entire group \mathbb{Z}_q^n , preventing structural shortcuts or biases.
- 317 • Non-Commutative Action: The non-commutative nature of $U \cdot t \pmod q$ disrupts
318 linear relationships, obfuscating the subgroup structure within H_{left} and making it
319 computationally infeasible to infer U or x directly.
- 320 • Combinatorial Complexity: The adversary must distinguish subgroup elements in
321 G_{ambient} , which involves an exponentially large search space without access to the
322 private transformations.

323 **Adversary’s Information-Theoretic Challenges:** In addition to the computational
324 hardness posed by NAHSP, the scheme introduces inherent information-theoretic barriers
325 that further obfuscate the relationships between inputs, outputs, and subgroup member-
326 ship:

- 327 • Uniformly Distributed Outputs: The transformations $t = A \cdot x \pmod q$ and $t' = U \cdot t$
328 $\pmod q$ ensure that the outputs are uniformly distributed over \mathbb{Z}_q^n . This uniformity
329 prevents adversaries from inferring structural dependencies or narrowing the search
330 space.
- 331 • Ambiguity of Subgroup Membership: Without access to the private matrix U or
332 secret vector x , adversaries cannot distinguish elements of the hidden subgroup N
333 from non-members within H_{left} , as the coset relationships are fully obscured.
- 334 • Exponential Pre-Image Set Sizes: The adversary must contend with an exponen-
335 tially large set of potential pre-images for any observed output, making it infeasible
336 to isolate the correct subgroup elements even under exhaustive search.

337 **Dual-Layered Security:**

338 The security of the scheme combines:

- 339 • Computational Hardness of NAHSP: The cryptographic guarantees are fundamen-
340 tally tied to the infeasibility of solving NAHSP, a problem resistant to both classical
341 and quantum adversaries.
- 342 • Information-Theoretic Obfuscation: The transformations induced by A , U , and x
343 ensure that observed outputs retain high entropy, preserving indistinguishability
344 across the full ambient group G_{ambient} .

345 **3 Problem Statement**

346 The objective of this cryptographic scheme is to secure the hidden subgroup N , ensuring
347 its structure remains concealed from adversaries. This is achieved by obfuscating the rela-
348 tionships between the public basis A , the hidden matrix U , and the secret vector x , while

349 leveraging a lossy mapping function $px(\cdot)$ to induce equivalence classes. The scheme em-
350 ploys dual-layered security mechanisms: computational hardness from the Non-Abelian
351 Hidden Subgroup Problem (NAHSP) and information-theoretic obfuscation from $px(\cdot)$.

352 Adversary's Knowledge

353 The adversary has access to:

- 354 • The public matrix A , which spans the ambient group G_{ambient} and defines H_{right} ,
- 355 • The obfuscated output t'' , resulting from applying the lossy mapping $px(\cdot)$ to ele-
356 ments of H_{left} .

357 Adversary's Limitations

358 The adversary does not have access to:

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

363 Equivocation of the Mapping Function $px(\cdot)$

364 The mapping function $px(\cdot)$ is a lossy transformation that projects elements of the am-
365 bient group G_{ambient} into equivalence classes. This mapping introduces significant ob-
366 fuscation, ensuring that subgroup membership cannot be determined feasibly without
367 knowledge of the private components.

368 Properties of $px(\cdot)$:

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

375 **Impact of Equivocation:** The lossy nature of $px(\cdot)$ exponentially increases the ad-
376 versary's search space by creating a many-to-one mapping:

- 377 • Pre-images of $px(\cdot)$ form equivalence classes that mask coset relationships within
378 H_{left} ,
- 379 • The adversary must contend with an exponentially large number of indistinguish-
380 able 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 $px(\cdot)$ is provably reduced through this lossy mapping. The surjectivity of $px(\cdot)$ 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 $k = 0$:

$$\begin{aligned} t_0 &\equiv B_0 \cdot x_0 \pmod{q}, \\ t'_0 &\equiv U_0 \cdot t_0 \pmod{q}, \\ t''_0 &\equiv px(t'_0 \circ J(C1_0)) \pmod{q}, \end{aligned}$$

For $k = 1$ to 5 :

$$\begin{aligned} t_k &\equiv B_k \cdot (x_k \circ t''_{k-1}) \pmod{q}, \\ t'_k &\equiv U_k \cdot t_k \pmod{q}, \\ t''_k &\equiv px(t'_k \circ J(C1_k)) \pmod{q}, \end{aligned}$$

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.

415 Verification and Forgery Mitigation

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

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

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

428 The verification equation is computed as follows:

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

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

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

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

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

431 A signature σ is valid if and only if:

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

432 Key Components:

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

444 **Forgery Mitigation:** Forgery is mitigated through the interaction of several mecha-
445 nisms:

- 446 • **Contextual Hash Constraints:** The hash constraints $C1$ and $C2$ bind the signature
447 and public key to specific contextual values, ensuring that signatures cannot be
448 reused or manipulated across different contexts.
- 449 • **Lossy Mapping Function and Probability of Forgery:** The lossy mapping function
450 $px(\cdot)$ obfuscates subgroup membership, making it computationally infeasible for
451 adversaries to generate valid signatures without access to the private keys. The
452 probability of a successful forgery at each level k is determined by the ratio of
453 equivalence class members m_k to the total number of equivalence classes n_k . For
454 each level, the adversary must generate a value that maps to the correct equivalence
455 class under $px(\cdot)$, resulting in a success probability of approximately $\frac{m_k}{n_k}$. Across K
456 levels, the cumulative probability of forging a valid signature is given by:

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

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

- 461 • **Fixed and Randomized Public Matrices:** The public matrix A is fixed and defines
462 the ambient group structure, while the signing matrix B is randomized and tied to
463 the message. This dynamic prevents adversaries from correlating multiple signa-
464 tures to infer structural relationships or exploit reuse.
- 465 • **Observed Entropy Constraint:** The observed entropy constraint ensures that edge-
466 case scenarios are effectively mitigated. Signature candidates are required to sta-
467 tistically adhere to near-randomness, aligning with approximately 1/10 of the com-
468 binatorial possibilities. This increases the difficulty of identifying predictable or
469 exploitable patterns, enhancing resilience to forgery.

470 **Soundness and Completeness:**

- 471 • **Completeness:** Any valid signature σ , generated using the correct private compo-
472 nents x and U , satisfies the verification equation.
- 473 • **Soundness:** Any invalid signature σ' , generated without access to the private compo-
474 nents, fails verification. This failure arises because σ' maps to incorrect equivalence
475 classes under $px(\cdot)$, and fails entropy checks for statistical validity.

476 *Related proofs of soundness, completeness, equivocation, and equivalence class ratio*
477 *impact will be presented as part of the formal proof of existential unforgeability under*
478 *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 $px(\cdot)$ 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 $px(\cdot)$ 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 $px(\cdot)$,
- 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 $px(\cdot)$, 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.

- 513 3. Verification through Structured Basis Transformations: Designing signature verifi-
514 cation as a proof of consistency through structured basis transformations. This ap-
515 proach ensures that transformations applied to the public key and signature align
516 under obfuscated subgroup relationships, preserving algebraic correctness while pre-
517 venting exploitation of subgroup structures.
- 518 4. Information-Theoretic Security via High-Entropy Mapping Functions: Introducing
519 a non-linear, many-to-one mapping function $px(\cdot)$ that compresses the ambient
520 space G_{ambient} into equivalence classes. The inherent high entropy of $px(\cdot)$'s outputs
521 enforces:
- 522 • Computational indistinguishability of equivalence class elements without ac-
523 cess to private keys x_k and U_k ,
 - 524 • Explicit rejection of low-entropy forgeries during verification, adding an entropy-
525 based security layer that complements computational hardness.
- 526 5. Novel Matrix-Based Cryptographic Framework: Developing a cryptographic plat-
527 form based on unrelated public/private matrix groups. The independence of public
528 matrices A , B , and private matrices U , x prevents structural exploitation, support-
529 ing efficient signature generation and verification.
- 530 6. First Practical Hybrid Information Theoretic and NAHSP-Based Construction: Con-
531 structing what we believe to be the first practical digital signature scheme explicitly
532 based on the NAHSP, using non-commutative matrix groups and leveraging both
533 information theoretic functions and chaining mechanisms to ensure robust security.

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

- 538 • **Communication Efficiency:** Preliminary parameters suggest competitive communi-
539 cation costs compared to RSA signatures and a significant reduction compared to
540 most lattice-based schemes.
- 541 • **Dual-Layered Security:** By combining information-theoretic indistinguishability with
542 computational hardness rooted in the NAHSP, the scheme provides a layered de-
543 fense against both classical and quantum adversaries.
- 544 • **Feasibility and Scalability:** The use of independent, unrelated public/private ma-
545 trix groups provides a scalable and tunable platform for balancing security and ef-
546 ficiency, with conservative parameter choices supporting incremental improvements
547 over time.

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

555 **Complexity Analysis** The mapping function $px(\cdot)$ partitions the group G into equiv-
 556 alence classes, obfuscating the subgroup structure of N and increasing the adversary’s
 557 difficulty in distinguishing elements. Unlike a group homomorphism, $px(\cdot)$ does not pre-
 558 serve group operations but ensures computational indistinguishability of elements within
 559 the same equivalence class. This indistinguishability amplifies the complexity of solving
 560 the problem by significantly increasing the effective solution space.

561 **Impact of $px(\cdot)$:** The hardness of the problem is tied directly to the pre-image count of
 562 $px(\cdot)$, which determines the size of equivalence classes and the adversarial search space.
 563 Specifically:

- 564 • The adversary must navigate all elements in $px^{-1}(g')$ for a given equivalence class
 565 g' , where $px^{-1}(g')$ contains all pre-images indistinguishable under $px(\cdot)$.
- 566 • The size of $px^{-1}(g')$ is determined by the partitioning of the ambient group G_{ambient}
 567 into equivalence classes via the mapping $px(\cdot)$. The hidden subgroup N and the
 568 transformations induced by x and U influence the structure of these partitions, but
 569 the pre-image membership size fundamentally scales with the size of G_{ambient} and
 570 configuration of $px(\cdot)$.
- 571 • The indistinguishability within equivalence classes ensures that structural relation-
 572 ships between elements of N and G are practically obscured, limiting adversarial
 573 insights.

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

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

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

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

585 **Security and Complexity Relationship:** The indistinguishability introduced by
 586 $px(\cdot)$ ensures that adversaries cannot efficiently distinguish elements of N within equiv-
 587 alence classes or between stages of the chaining mechanism. By explicitly tying the
 588 complexity to the pre-image count $|px^{-1}(g')|$, the scheme achieves:

- 589 • Scalable Hardness: The size of equivalence classes and the number of chained in-
 590 stances can be tuned to balance efficiency and security.
- 591 • Resistance to Structural Attacks: The obfuscation introduced by $px(\cdot)$ disrupts struc-
 592 tural relationships, ensuring that subgroup recovery requires infeasible computa-
 593 tional 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 $px(\cdot)$ 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 $px(\cdot)$ 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 Notation and Definitions

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.

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

6.1 Group Structure and Properties

Ambient Group G_{ambient} : The ambient group G_{ambient} is defined as:

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

the group of invertible $n \times n$ matrices over \mathbb{Z}_q , where the group operation is matrix multiplication modulo q . This group is non-abelian and serves as the foundational structure for constructing G .

Working Group G : The working group G is constructed as a semidirect product:

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

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 .

The automorphism action of H_{right} on H_{left} ensures that G is non-abelian:

$$h_R \cdot h_L \cdot h_R^{-1} = \phi_{h_R}(h_L), \quad h_R \in H_{\text{right}}, \quad h_L \in H_{\text{left}},$$

where ϕ_{h_R} is an automorphism of H_{left} .

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 :

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

6.2 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 $f : G \rightarrow S$, 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 .

6.3 Oracle Construction and Reduction to NAHSP

Oracle Function f : Define the oracle function $f : G \rightarrow S$ 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.

6.4 Hardness of Subgroup Recovery

- **Non-Abelian Structure:** The semidirect 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 . \square

6.5 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 $f : G \rightarrow S$, 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 $f : G \rightarrow S$ 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.

- 707 • **Efficiency:** The reduction \mathcal{R} invokes \mathcal{A} once and performs polynomial-time group
708 operations for verification. Therefore, the computational overhead of \mathcal{R} is polyno-
709 mial in the size of G and bounded by the runtime of \mathcal{A} .
- 710 • **Adversarial Advantage:** Suppose \mathcal{A} has a non-negligible advantage ϵ in recover-
711 ing N . Then, \mathcal{R} achieves the same advantage in solving NAHSP:

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

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

714 Hardness of Subgroup Recovery.

- 715 • **Classical Adversaries:** Classical algorithms would need to enumerate cosets of N ,
716 which is computationally infeasible due to the exponential size of G . Additionally,
717 the non-abelian structure of G lacks the necessary algebraic properties that allow
718 for efficient subgroup identification, preventing the use of techniques such as brute-
719 force search or combinatorial optimizations.
- 720 • **Quantum Adversaries:** Quantum algorithms, including those leveraging the
721 Quantum Fourier Transform (QFT), struggle with G 's non-abelian structure. The
722 automorphism action of H_{right} on H_{left} disrupts the periodicity and coherence es-
723 sential for QFT-based subgroup recovery. As a result, these quantum techniques
724 fail to efficiently exploit the hidden subgroup structure in G , ensuring resistance
725 against known quantum attacks.

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

730 7 Equivocation and Indistinguishability from $px(\cdot)$

731 7.1 Equivocation Function $px(\cdot)$

732 The equivocation function $px(\cdot)$ is a deterministic mapping that compresses an input
733 element into an equivalence class represented by the output. It is a lossy, surjective,
734 many-to-one compression function that reduces real entropy while maintaining high ob-
735 served entropy in the output. The mutual entropy between the input and output is
736 distributed across indistinguishable equivalence classes, ensuring computational imprac-
737 ticality in enumerating all potential valid inputs from a given output. While the correct
738 input is guaranteed to exist within the equivalence class, no heuristic information can
739 differentiate it from other valid pre-images. Note that we use level III parameters for this
740 section in general.

741 The $px(\cdot)$ function operates as follows:

742 **1. Inverse NTT Transform:** The input element, initially represented in the NTT
743 domain over a field $q = 257$, is transformed back to the input domain via the appropriate
744 inverse NTT, based on n .

745 **2. Forward NTT Transform to $q = 1283$:** A forward NTT is performed using
 746 parameters $q = 1283, \omega = 3$. Note that only addition is performed using this field, so
 747 using $\omega = 3$ only impacts computational performance. This step expands the coefficient
 748 range by approximately a factor of 5, mapping them to values between 0 and 1282.

749 **3. Scaling and Ambiguity Introduction:** The element is added to itself element-
 750 wise and reduced modulo 1283. This scaling operation is repeated four times, introducing
 751 additional ambiguity at each step. For each coefficient, there are two possible pre-image
 752 states—either it rolled over or it did not. This process creates an internal diffusion factor
 753 of 2^{4n} .

754 **4. Uniform Input Distribution:** The input element is derived from:

- 755 • A uniformly random, full-rank hidden matrix U ,
- 756 • A uniquely randomized, full-rank public matrix B ,
- 757 • A uniformly random secret element x .

758 These components ensure the input uniformly spans the ambient modular space \mathbb{Z}_q^n ,
 759 directly supporting the uniformity produced by the diffusion factor 2^{4n} . **5. Inverse and**
 760 **Forward NTT Transforms:** The diffused element undergoes:

- 761 • An inverse NTT transform with parameters $q = 1283, \omega = 3$,
- 762 • A forward NTT transform with parameters $q = 257, \omega = \{81, 9, 3\}$, for on $n =$
 763 $64, n = 128, n = 256$ respectively.

764 This step maps the element back to a smaller field while preserving ambiguity.

765 Each coefficient in the $q = 1283$ field has approximately 5 pre-image coefficients in
 766 the $q = 257$ field. With n coefficients, the total number of pre-images is:

$$5^n.$$

767 For $n = 128$, this results in 2^{297} pre-images distributed across 2^{727} equivalence classes
 768 (assuming uniform partitioning). Each equivalence class contains 2^{297} elements, making
 769 it computationally infeasible for an adversary to enumerate all valid pre-images for a
 770 given output. The core of our information-theoretic security lies in the impracticality of
 771 inverting $px(\cdot)$. This security is based on the following principles:

772 Information-Theoretic Security

773 The core of our information-theoretic security lies in the impracticality of inverting $px(\cdot)$.
 774 This security is based on the following principles:

- 775 • Diffusion and Avalanche Effect: The diffusion factor 2^{4n} ensures that small changes
 776 in the input lead to significant and widespread changes in the output, making it
 777 difficult to trace back to the original input.
- 778 • Pre-Image Resistance: Mapping to approximately 5^n pre-images distributed across
 779 2^{727} equivalence classes for $n = 128$ ensures that each output corresponds to a large
 780 and computationally infeasible set of inputs.

- 781 • NTT Transformation Security: The use of NTTs with carefully chosen parameters
782 introduces additional complexity, leveraging the hardness of problems related to
783 discrete transforms over finite fields.
- 784 • Assumptions on Computational Resources: The security guarantees assume that
785 adversaries do not possess exponential computational resources to perform exhaus-
786 tive searches within the equivalence classes.

787 In summary, while the $px(\cdot)$ function allows enumeration of all pre-images theoreti-
788 cally, the computational and combinatorial requirements make this infeasible in practice.
789 This foundational design ensures the cryptographic strength of the equivocation function
790 in supporting secure operations.

791 7.2 Analysis of $px()$

792 This section rigorously establishes the security properties of the mapping function $px(\cdot)$,
793 a core cryptographic primitive in this system. Through formal proofs, we demonstrate its
794 resilience against adversarial attacks, its statistical uniformity, highlighting its robustness
795 in both classical and quantum computational models.

796 It should be noted that selecting specific parameters for our $px()$ function requires
797 consideration of modular overlap. The foundation of our function is based on additive
798 diffusion of projected elements from $q_0 = 257$ to a larger prime field, such as $q_1 = 1283$.
799 As it is mathematically impossible for two prime numbers to divide cleanly, we seek to get
800 as close as possible. For example, a $q = 1285 == 5 * 257$, leaving no modular remainder
801 during mapping, but 1285 is not a prime number. Our choice of $q_1 = 1283$ in this case
802 leaves us 1.17% overlapping values during inversion. Restated, out of 257 elements 254
803 will have 5 possible pre-image coefficients that map back, with the remainder having 4.
804 The implication being that in reality, not all equivalence classes hold exactly the same
805 number of pre-images and there is a slight deviation related to the 1.17% modular overlap.
806 For simplicity, as the majority in this case have 5 valid pre-image coefficients, we will use
807 5, vs a more accurate $\approx 4.99 - 5.01$ value.

808 7.3 Probability of Random Forgery

809 The information-theoretic barrier we create is not infinite, but it presents an intractably
810 large solution space for an adversary, as we will show below.

Table 2: Preimage and Equivalence Class Analysis for $px()$

Dimension (n)	Ambient Space (2^n)	Preimages (2^b)	Equivalence Classes (2^e)
64	2^{512}	2^{149}	2^{363}
128	2^{1024}	2^{297}	2^{727}
256	2^{2048}	2^{640}	2^{1408}

811 These numbers yield the following probability ratios:

$$\begin{aligned}
 \text{For } n = 64 : & \quad 1.844674407370955 \times 10^{-45}, \\
 \text{For } n = 128 : & \quad 3.402823669209385 \times 10^{-90}, \\
 \text{For } n = 256 : & \quad 1.157920892373162 \times 10^{-179}.
 \end{aligned}$$

812 This represents the probability for each n , that a randomly selected pre-image will
 813 map to a selected equivalence class. These numbers are useful for proving probability of
 814 random signature forgery, but to put them into context we will relate them to winning
 815 the PowerBall lottery which has a probability of 1 in 292,201,338. The only constraint is
 816 that you can only buy one ticket per jackpot.

n	Probability	Powerball Wins	Probability (Chaining)	Powerball Wins (Chaining)
64	1.84×10^{-45}	13	1.34×10^{-358}	350
128	3.40×10^{-90}	26	1.55×10^{-537}	529
256	1.15×10^{-180}	53	1.80×10^{-720}	711

817 Assuming these calculations are correct, without chaining the odds of guessing a valid
 818 signature for $n = 64$ are \approx the same as winning PowerBall 13 times in a row. As level
 819 I has $n = 64$ and 8 chained instances, the success probability, based on buying one
 820 ticket per jackpot, jumps to winning PowerBall 350 times in a row. Guessing a valid
 821 signature at level V where $n = 256$ with 4 chained instances is as likely as winning 711
 822 consecutive jackpots. As this is technically not impossible, similar to how the proposed
 823 scheme doesn't guarantee perfect secrecy, it could in theory happen. It's just unlikely.

824 7.3.1 Brute Force Storage Requirements

825 In the context of attempting to mount a brute force attack, we consider the minimum stor-
 826 age requirements to guarantee a successful outcome. Each element requires $\{64, 128, 256\}$
 827 bytes of storage and each observed output t'' maps to a large number of valid preimages.
 828 There also exists a lower bound of observations required to guarantee enough correct t'
 829 outputs such that the correct set of equations that can be defined and solved to recover
 830 both U and x . The storage required for each security level is listed below, in brontobytes:

Table 3: Preimage Storage and Observation Scaling for $px()$

n	Preimages (2^b)	1 Observation (BB)	Min Observations (BB)
64	2^{149}	2^{65}	2^{72}
128	2^{297}	2^{214}	2^{221}
256	2^{640}	2^{648}	2^{656}

831 Brontobytes and Exabytes: Units for Massive Data Volumes

832 The **brontobyte** (BB) is a theoretical unit of data storage equivalent to:

$$1 \text{ BB} = 2^{90} \text{ bytes} = 1,024 \text{ yottabytes.}$$

833 It represents a continuation of the binary progression of data storage units. To provide
 834 perspective:

- 835 • 1 BB = 1,024 YB (yottabytes),
- 836 • 1 YB = 1,024 ZB (zettabytes),
- 837 • 1 ZB = 1,024 EB (exabytes).

838 By comparison:

839 • 1 EB = 2^{60} bytes, or roughly 1 billion gigabytes.

840 • A single brontobyte (2^{90}) is 1,099,511,627,776 EB, far exceeding the total amount
841 of data stored globally.

842 Relevance to Information-Theoretic Security

843 In this context, brontobyte-scale data is required for analyzing the feasibility of brute-force
844 attacks and quantifying the strength of information-theoretic guarantees. For example:

845 • Storing all preimages for even a single observation exceeds brontobyte (BB) levels
846 with increasing dimensionality.

847 • As of 2024, the global installed base of data storage capacity is projected to be
848 approximately 11.23 zettabytes (ZB), equivalent to 11.23 trillion terabytes (TB) or
849 0.00001071 brontobytes (BB).

850 • For minimum required observations ($n = 64, 65$ observations), this storage require-
851 ment guarantees an adversary faces infeasible data handling and computational
852 costs.

853 • Brute force storage requirements are infeasible, and the probability of simply guess-
854 ing a valid answer is negligible, exactly what one would expect from practical in-
855 formation theoretic security guarantees.

856 7.4 Equivalence Class Invariance of $px(\cdot)$

857 **Lemma 1** (Equivalence Class Invariance of $px(\cdot)$). *The mapping function $px : G \rightarrow \mathcal{Y}$*
858 *partitions the group G into equivalence classes defined by the subgroup N . Specifically:*

$$px(g_1) = px(g_2) \iff g_1 \sim g_2 \quad \forall g_1, g_2 \in G,$$

859 where \sim is the equivalence relation:

$$g_1 \sim g_2 \iff g_1 \cdot n = g_2 \text{ for some } n \in N.$$

860 *Proof. Objective:* To show that $px(\cdot)$ is invariant over equivalence classes induced by
861 N .

862 **Definition of Equivalence Classes:** The equivalence class of an element $g \in G$ with
863 respect to N is:

$$[g]_N = \{g \cdot n \mid n \in N\}.$$

864 These classes partition G into disjoint subsets such that:

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

865 **Invariance of $px(\cdot)$:** The mapping function $px(\cdot)$ satisfies the following invariance prop-
 866 erty:

$$px(g \cdot n) = px(g) \quad \forall g \in G, n \in N.$$

867 This property implies that $px(\cdot)$ maps all elements of an equivalence class $[g]_N$ to the
 868 same value in \mathcal{Y} .

869 **Proof of Equivalence:** To prove $px(g_1) = px(g_2) \iff g_1 \sim g_2$, we consider both
 870 directions:

871 • **Forward Direction (\Rightarrow):** Assume $px(g_1) = px(g_2)$. By the invariance property
 872 of $px(\cdot)$, this implies:

$$px(g_1 \cdot n_1) = px(g_2 \cdot n_2) \quad \text{for some } n_1, n_2 \in N.$$

873 Since $px(\cdot)$ is consistent across equivalence classes, it follows that:

$$g_1 \cdot n_1 = g_2 \cdot n_2 \quad \text{for some } n_1, n_2 \in N,$$

874 which implies $g_1 \sim g_2$.

875 • **Backward Direction (\Leftarrow):** Assume $g_1 \sim g_2$, i.e., $g_1 \cdot n = g_2$ for some $n \in N$. By
 876 the invariance property of $px(\cdot)$:

$$px(g_2) = px(g_1 \cdot n) = px(g_1).$$

877 Therefore, $px(g_1) = px(g_2)$.

878 **Conclusion:** The function $px(\cdot)$ is invariant across equivalence classes induced by N .
 879 This ensures that elements within the same equivalence class are indistinguishable under
 880 $px(\cdot)$. □

881 7.5 Enhanced Hiding of $px(\cdot)$

882 **Lemma 2** (Ambiguity and Obfuscation in $px(\cdot)$). *The mapping function $px(\cdot)$ is a*
 883 *many-to-one function that increases the entropy of its outputs and introduces obfus-*
 884 *cation through the hidden matrix U . Specifically, $px(\cdot)$ creates an ambiguous search space*
 885 *where recovering the original inputs is computationally infeasible.*

886 *Proof. Objective:* To show that $px(\cdot)$ creates ambiguity by mapping multiple distinct
 887 inputs to the same output and obfuscates structural relationships through the hidden
 888 matrix U .

889 **Many-to-One Mapping:** The mapping $px(\cdot)$ compresses the input space, assigning
 890 multiple distinct inputs to the same output:

$$|px^{-1}(y)| \geq 2^k, \quad \text{where } k \text{ depends on system parameters.}$$

891 This many-to-one nature ensures that adversaries cannot uniquely identify an input from
 892 an output.

893 **Obfuscation via U :** The hidden matrix U transforms the input as:

$$t' = U \cdot t \pmod{q}.$$

894 Without knowledge of U , inverting this transformation is computationally infeasible. The
895 randomness of U ensures that no exploitable dependencies exist between the input t and
896 its transformation t' .

897 **Combined Effect:** The combined properties of $px(\cdot)$ and U ensure:

- 898 • **Ambiguity:** Each output corresponds to a large equivalence class of indistinguish-
899 able inputs, creating an inflated search space.
- 900 • **Obfuscation:** Structural relationships between inputs are disrupted, preventing
901 adversaries from reconstructing t without explicit knowledge of U .

902 **Security Implications:** Recovering the original input t for a given output $px(t)$ re-
903 quires exhaustive search through all possible pre-images. The exponential size of the
904 search space and the disruption of algebraic structures ensure that this task is computa-
905 tionally infeasible.

906 **Conclusion:** The mapping $px(\cdot)$, combined with the obfuscation introduced by U , cre-
907 ates a high-entropy, ambiguous output space. These properties ensure the security of the
908 mapping against adversarial recovery of original inputs. \square

909 7.6 Indistinguishability and Statistical Uniformity of $px(\cdot)$ Out- 910 puts

911 Note that we use Level III parameters for this section, $n = 128, q = 257$.

912 **Lemma 3** (Indistinguishability of $px(\cdot)$ Outputs). *The outputs of the mapping function*
913 *$px(\cdot): \mathbb{Z}_{257}^{128} \rightarrow \mathbb{Z}_{257}^{128}$ are computationally indistinguishable from uniformly random vec-*
914 *tors in \mathbb{Z}_{257}^{128} , given no access to the secret input x , intermediate values t' , or the hidden*
915 *transformation parameters.*

916 *Proof. Objective:* To prove that the outputs $px(t')$ are computationally indistinguish-
917 able from uniformly random vectors, assuming adversaries lack access to the secret input
918 x or intermediate values.

919 **Step 1: High Entropy and Uniform Mixing** The mapping $px(\cdot)$ introduces high
920 entropy and uniform mixing through sequential transformations:

- 921 • **Initial Transformation (NTT):** The input vector t is transformed into $t' =$
922 $\text{NTT}_{1283}(t)$, dispersing coefficients of t across the frequency domain. This ensures
923 pseudorandom spreading of t' over \mathbb{Z}_{1283}^{128} .
- 924 • **Additive Mixing:** The operation $t'' = t' + 2 \cdot t'$ introduces further uniformity,
925 erasing any residual structure in t' .
- 926 • **Inverse Transformation:** The inverse NTT, $z = \text{INV_NTT}_{1283}(t'')$, preserves
927 high entropy and disperses any remaining correlations across coefficients in \mathbb{Z}_{1283}^{128} .
- 928 • **Final Projection:** The mapping $px(t') = \text{NTT}_{257}(z)$ reduces z modulo 257, en-
929 suring outputs are uniformly distributed in \mathbb{Z}_{257}^{128} and removing residual patterns.

930 **Step 2: Compression and Ambiguity** The function $px(\cdot)$ compresses \mathbb{Z}_{1283}^{128} into \mathbb{Z}_{257}^{128} ,
 931 introducing a many-to-one mapping. Each output $y \in \mathbb{Z}_{257}^{128}$ corresponds to approximately
 932 2^{297} indistinguishable pre-images (for level III):

$$|px^{-1}(y)| \approx 2^{297}.$$

933 This compression ensures that adversaries observing $px(t')$ cannot deduce a unique input
 934 t , significantly inflating the effective search space.

935 **Step 3: Statistical Uniformity** The modular reductions and transformations in $px(\cdot)$
 936 ensure that outputs pass standard randomness tests:

- 937 • **Entropy Preservation:** High entropy at intermediate states t' and z ensures no
 938 statistical patterns remain.
- 939 • **Empirical Validation:** Statistical tests (e.g., NIST randomness suite) confirm
 940 that $px(t')$ outputs are indistinguishable from uniformly random elements in \mathbb{Z}_{257}^{128} .

941 **Step 4: Entropy and Adversarial Uncertainty** The lossy nature of $px(\cdot)$ guaran-
 942 tees high apparent entropy for adversaries:

- 943 • **Min-Entropy (H_∞):** Assuming uniform distribution over \mathbb{Z}_{257}^{128} , the min-entropy
 944 of $px(t')$ is:

$$H_\infty(px(t')) = \log_2(|\mathbb{Z}_{257}^{128}|) = 727 \text{ bits.}$$

- 945 • **Conditional Entropy ($H(t' | px(t'))$):** Given $px(t')$, the adversary faces residual
 946 uncertainty about t' :

$$H(t' | px(t')) = 1024 - 727 = 297 \text{ bits.}$$

947 This indicates that each output corresponds to 2^{297} indistinguishable pre-images,
 948 obfuscating input-output relationships.

- 949 • **False Entropy Perception:** From the adversary's perspective, $px(t')$ appears to
 950 have full entropy $H(px(t'))$, as outputs are indistinguishable from uniform distri-
 951 butions.

952 **Step 5: Computational Indistinguishability** For any efficient adversary \mathcal{A} , distin-
 953 guishing $px(t')$ from a uniformly random vector $r \in \mathbb{Z}_{257}^{128}$ is computationally infeasible:

$$|\Pr[\mathcal{A}(px(t')) = 1] - \Pr[\mathcal{A}(r) = 1]| \leq \text{negl}(n),$$

954 where $\text{negl}(n)$ is a negligible function of the security parameter n . The modular reduc-
 955 tions and many-to-one compression ensure that adversaries cannot exploit patterns to
 956 distinguish $px(t')$ from random vectors.

957 **Conclusion** The mapping $px(\cdot)$ ensures high entropy, statistical uniformity, and com-
 958 putational indistinguishability. These properties collectively enhance its cryptographic
 959 strength, making the outputs indistinguishable from uniformly random vectors and robust
 960 against adversarial analysis. □

961 7.7 Collision and Pre-Image Resistance of $px(\cdot)$

962 **Lemma 4** (Collision and Pre-Image Resistance of $px(\cdot)$). *The mapping function $px(\cdot): \mathbb{Z}_{257}^{128} \rightarrow$*
 963 *\mathbb{Z}_{257}^{128} satisfies:*

- 964 1. **Collision Resistance:** *Finding distinct inputs $t_1, t_2 \in \mathbb{Z}_{257}^{128}$ such that $px(t_1) =$*
 965 *$px(t_2)$ is computationally infeasible.*
- 966 2. **Pre-Image Resistance:** *Given $y \in \mathbb{Z}_{257}^{128}$, finding any $t \in \mathbb{Z}_{257}^{128}$ such that $px(t) = y$*
 967 *is computationally infeasible.*

968 *These properties hold under standard cryptographic assumptions.*

969 *Proof. Objective:* To demonstrate that $px(\cdot)$ is resistant to both collision and pre-image
 970 attacks by analyzing its structure, randomness, and computational complexity.

971 **Structure of $px(\cdot)$** The mapping $px(\cdot)$ consists of the following steps:

- 972 1. $t' = \text{NTT}_{1283}(t)$, where $t = A \cdot x \pmod{257}$ and A is a random public matrix.
- 973 2. $t'' = t' + 3 \cdot t'$, introducing additive mixing in \mathbb{Z}_{1283}^{128} .
- 974 3. $z = \text{INV_NTT}_{1283}(t'')$, returning to the time domain modulo 1283.
- 975 4. $px(t) = \text{NTT}_{257}(z)$, projecting the result into \mathbb{Z}_{257}^{128} .

976 These transformations ensure randomness, mixing, and compression, making $px(\cdot)$ resis-
 977 tant to both collision and inversion attempts.

978 **Collision Resistance Analysis** For $px(t_1) = px(t_2)$ to hold, it must be true that
 979 $z_1 = z_2$, since NTT_{257} is invertible. This implies:

$$\text{INV_NTT}_{1283}(t''_1) = \text{INV_NTT}_{1283}(t''_2).$$

980 Given that:

$$t''_1 = t'_1 + t'_1, \quad t''_2 = t'_2 + t'_2,$$

981 distinct $t'_1 \neq t'_2$ result in distinct $t''_1 \neq t''_2$ due to additive mixing. A collision would require:

$$\text{NTT}_{1283}(t_1) = \text{NTT}_{1283}(t_2),$$

982 which is unlikely given the pseudorandom nature of A . The probability of such random
 983 collisions is bounded by:

$$P_{\text{collision}} \leq \frac{q^2}{2 \cdot 257^{128}},$$

984 where q is the number of adversarial queries. Since 257^{128} is astronomically large, $px(\cdot)$
 985 is collision-resistant.

986 **Pre-Image Resistance Analysis** To invert $px(t) = y$, an adversary must reverse
987 multiple transformations:

- 988 • Recover z from $y = \text{NTT}_{257}(z)$, which is infeasible without knowledge of t or inter-
989 mediate states.
- 990 • Reverse $z = \text{INV_NTT}_{1283}(t'')$ to find t'' , where $t'' = t' + 3 \cdot t'$. Additive mixing
991 obscures linear relationships in t' .
- 992 • Solve $t = A \cdot x \pmod{257}$ from $t' = \text{NTT}_{1283}(t)$. Without knowledge of x , this is
993 computationally infeasible due to the pseudorandomness of A .

994 Additionally, the lossy nature of $px(\cdot)$ ensures:

$$|px^{-1}(y)| \approx 2^{297},$$

995 making it computationally infeasible to identify a unique pre-image among 2^{297} candi-
996 dates.

997 **Conclusion** The structural properties of $px(\cdot)$, including pseudorandom transforma-
998 tions, additive mixing, and modular reductions, ensure resistance to both collision and
999 pre-image attacks. These properties make $px(\cdot)$ secure under standard cryptographic
1000 assumptions against both classical and quantum adversaries. \square

1001 7.8 Avalanche Effect in $px(\cdot)$

1002 **Lemma 5** (Avalanche Effect of $px(\cdot)$). *For any inputs $t_1, t_2 \in \mathbb{Z}_{257}^{128}$ differing by a single*
1003 *bit, the outputs $px(t_1)$ and $px(t_2)$ are computationally indistinguishable from independent,*
1004 *uniformly random vectors in \mathbb{Z}_{257}^{128} .*

1005 *Proof. Objective:* To show that a small change in t propagates unpredictably through
1006 $px(\cdot)$, ensuring significant and uncorrelated differences in the outputs.

1007 **Step 1: Propagation Through NTT Transformations** The input t undergoes the
1008 transformation $t' = \text{NTT}_{1283}(t)$. Due to the properties of the NTT:

- 1009 • Each coefficient of t' depends on all coefficients of t .
- 1010 • A single-bit change in t affects every coefficient of t' due to the frequency-domain
1011 dispersion.

1012 This ensures that the effect of a single-bit change is amplified across the intermediate
1013 state t' .

1014 **Step 2: Additive Mixing** The operation $t'' = t' + 2 \cdot t'$ introduces further mixing:

- 1015 • The additive mixing operation is performed modulo 1283, ensuring that changes
1016 propagate unpredictably due to modular wraparound.
- 1017 • Any change in t' affects all coefficients of t'' .

1018 **Step 3: Inverse Transformation** The inverse NTT $z = \text{INV_NTT}_{1283}(t'')$ maps the
 1019 mixed state back to the time domain. This operation:

- 1020 • Preserves the non-linear dependencies introduced by additive mixing.
- 1021 • Further disperses the effects of the initial change across all coefficients of z .

1022 **Step 4: Final Projection** The projection to \mathbb{Z}_{257}^{128} via NTT_{257} ensures that the output
 1023 $px(t)$ reflects the amplified changes from earlier stages. Specifically:

- 1024 • Modular reduction ensures that even small differences in z produce large, unpre-
 1025 dictable differences in $px(t)$.
- 1026 • The NTT modulo 257 further spreads any changes across all coefficients.

1027 **Step 5: Statistical Indistinguishability** For any t_1, t_2 differing by a single bit, the
 1028 outputs $px(t_1)$ and $px(t_2)$ satisfy:

$$\Pr[\mathcal{A}(px(t_1)) = 1] - \Pr[\mathcal{A}(px(t_2)) = 1] \leq \text{negl}(n),$$

1029 where \mathcal{A} is any efficient adversary and $\text{negl}(n)$ is a negligible function of the security
 1030 parameter n .

1031 **Conclusion** The transformations within $px(\cdot)$ amplify any small changes in t , ensuring
 1032 that $px(t_1)$ and $px(t_2)$ are computationally indistinguishable from independent, uniformly
 1033 random vectors. This establishes the avalanche effect for $px(\cdot)$. \square

1034 7.9 Adaptive Security of $px(\cdot)$

1035 **Lemma 6** (Adaptive Security of $px(\cdot)$). *For any adversary \mathcal{A} making up to q adaptive*
 1036 *queries to $px(\cdot)$, the outputs of $px(\cdot)$ remain indistinguishable from independent, uniformly*
 1037 *random vectors in \mathbb{Z}_{257}^{128} , given the randomized matrix A is independently regenerated for*
 1038 *each operation.*

1039 *Proof. Objective:* To prove that $px(\cdot)$ maintains its security properties against adver-
 1040 saries making multiple adaptive queries.

1041 **Step 1: Randomization of A** The matrix A is independently regenerated for each
 1042 invocation of $px(\cdot)$. This ensures that:

- 1043 • Outputs $px(t)$ from different invocations are uncorrelated.
- 1044 • An adversary cannot infer patterns or dependencies between outputs from different
 1045 queries.

1046 **Step 2: Independence of Transformations** Each invocation of $px(\cdot)$ is independent
 1047 due to the randomized A . Specifically:

- 1048 • The NTT transformations NTT_{1283} and NTT_{257} depend on A , ensuring fresh ran-
 1049 domness for each query.
- 1050 • Additive mixing and modular reductions are independent for each invocation, fur-
 1051 ther decoupling the outputs.

1052 **Step 3: Indistinguishability Under Adaptive Queries** For any q adaptive queries
 1053 t_1, t_2, \dots, t_q , the corresponding outputs $px(t_1), px(t_2), \dots, px(t_q)$ are indistinguishable
 1054 from independent, uniformly random vectors. Formally:

$$\Delta = |\Pr[\mathcal{A}(px(t_1), \dots, px(t_q)) = 1] - \Pr[\mathcal{A}(r_1, \dots, r_q) = 1]| \leq \text{negl}(n),$$

1055 where r_1, \dots, r_q are independent, uniformly random vectors in \mathbb{Z}_{257}^{128} .

1056 **Step 4: Resilience to Query Correlations** Even if \mathcal{A} chooses t_1, t_2, \dots, t_q adaptively,
 1057 the randomized A ensures that:

- 1058 • Outputs $px(t_i)$ are uncorrelated.
- 1059 • Knowledge of $px(t_i)$ does not provide any advantage in predicting $px(t_{i+1})$.

1060 **Conclusion** The independence of A across queries ensures that $px(\cdot)$ is secure against
 1061 adaptive adversaries, maintaining indistinguishability and unpredictability under multiple
 1062 queries. \square

1063 7.10 Adversarial Complexity and Relation to NAHSP

1064 **Lemma 7** (Adversarial Complexity of Pre-Image Recovery). *Recovering the valid pre-*
 1065 *image of $px(t')$ requires brute-forcing all 2^{297} indistinguishable pre-images and testing*
 1066 *each against the cryptographic construction. This task is computationally infeasible under*
 1067 *both classical and quantum adversarial models, as it reduces to solving a combinatorial*
 1068 *subgroup recovery problem tied to NAHSP.*

1069 *Proof. Objective:* To show that recovering the valid pre-image of $px(t')$ is computa-
 1070 tionally infeasible due to the obfuscation introduced by $px(\cdot)$ and its connection to the
 1071 Non-Abelian Hidden Subgroup Problem (NAHSP).

1072 **Step 1: Compression and Pre-Image Ambiguity** The mapping function $px(\cdot)$
 1073 transforms inputs $t \in \mathbb{Z}_{1283}^{128}$ to outputs $px(t') \in \mathbb{Z}_{257}^{128}$, with a compression ratio of approx-
 1074 imately 2^{297} -to-1:

$$R_{\text{compression}} = \frac{|\mathbb{Z}_{1283}^{128}|}{|\mathbb{Z}_{257}^{128}|} = 2^{297}.$$

1075 For a given output $px(t') = y$, the adversary faces approximately 2^{297} indistinguishable
 1076 pre-images $t_1, t_2, \dots, t_{2^{297}}$. Among these, only one pre-image corresponds to the correct
 1077 subgroup N .

1078 **Step 2: Valid Pre-Image and Subgroup Recovery** The valid pre-image satisfies
 1079 the transformation:

$$t' = U \cdot t \pmod{q}, \quad t = A \cdot x \pmod{q},$$

1080 where:

- 1081 • $A \in \mathbb{Z}_{257}^{128 \times 128}$ is a public, randomized, full-rank matrix.
- 1082 • $U \in \mathbb{Z}_{257}^{128 \times 128}$ is a secret, dense, full-rank matrix defining the subgroup N .

1083 Recovering this valid pre-image is equivalent to solving the subgroup recovery problem for
 1084 N in $G = H_{\text{left}} \rtimes H_{\text{right}}$, where G is the semidirect product of $H_{\text{left}} = \langle U \rangle$ and $H_{\text{right}} = \langle A \rangle$.

1085 **Step 3: Combinatorial Search Space** Without knowledge of x or U , the adversary
1086 must:

- 1087 • Enumerate all 2^{297} indistinguishable pre-images t_i for the given output $px(t') = y$.
- 1088 • Test each pre-image against the cryptographic construction to determine whether
1089 it satisfies the subgroup structure defined by A and U .

1090 This brute-force search involves solving a system of obfuscated equations for each candi-
1091 date t_i , including:

- 1092 • Modular reductions in \mathbb{Z}_{1283} and \mathbb{Z}_{257} ,
- 1093 • Non-linear dependencies introduced by NTT transformations and additive mixing.

1094 The total complexity scales as:

$$O(2^{297}),$$

1095 since each pre-image requires testing against the subgroup structure.

1096 **Step 4: Reduction to NAHSP** The task of identifying the valid pre-image reduces
1097 to solving the Non-Abelian Hidden Subgroup Problem (NAHSP) for G :

- 1098 • The hidden subgroup N is defined by $H_{\text{left}} = \langle U \rangle$.
- 1099 • The cosets of N in G correspond to equivalence classes of inputs under $px(\cdot)$.

1100 Solving the NAHSP involves identifying the subgroup N from its coset structure, which
1101 is computationally hard for non-abelian groups like G . The obfuscation introduced by
1102 $px(\cdot)$ ensures that:

$$\Pr[\text{Adversary recovers } N] \leq \frac{1}{|px^{-1}(y)|} = \frac{1}{2^{297}}.$$

1103 **Step 5: Resistance to Quantum Speedup** Quantum algorithms like Grover's pro-
1104 vide no advantage because:

- 1105 • The search space is structured around the subgroup recovery problem for N , which
1106 involves combinatorial dependencies between pre-images.
- 1107 • NAHSP inherently disrupts the coherence and periodicity necessary for quantum
1108 algorithms to achieve efficient speedups.
- 1109 • The adversary must brute-force permutations of obfuscated equations, which cannot
1110 be accelerated by Grover's algorithm.

1111 **Step 6: Formal Complexity Analysis** The total complexity of recovering the valid
1112 pre-image can be summarized as:

- 1113 • **Classical Complexity:** $O(2^{297})$, due to the need to brute-force all indistinguish-
1114 able pre-images.
- 1115 • **Quantum Complexity:** $O(2^{297})$, as quantum algorithms provide no advantage
1116 for structured subgroup recovery problems.

1117 **Conclusion** Recovering the valid pre-image of $px(t')$ reduces to solving the NAHSP
 1118 for $G = H_{\text{left}} \rtimes H_{\text{right}}$. The compression introduced by $px(\cdot)$, combined with the obfusca-
 1119 tion from modular reductions, NTTs, and subgroup structures, ensures that this task is
 1120 computationally infeasible for both classical and quantum adversaries. \square

1121 7.11 Quantum Resistance of $px(\cdot)$

1122 The mapping $px(\cdot)$ achieves quantum resistance by introducing high entropy, compression,
 1123 and structural obfuscation, effectively neutralizing known quantum algorithmic advan-
 1124 tages. Key disruptions include:

1125 • Quantum Fourier Transform (QFT):

- 1126 – $px(\cdot)$ collapses cosets of G into indistinguishable equivalence classes, removing
 1127 the periodic eigenstate structures required for QFT-based solvers.
- 1128 – The non-abelian properties of G disrupt coherence and prevent the exploitation
 1129 of group symmetries, negating QFT efficiency.

1130 • Grover’s Search:

- 1131 – Compression and indistinguishability inflate the effective search space, coun-
 1132 teracting Grover’s quadratic speedup by increasing the adversary’s uncertainty
 1133 over 2^{297} indistinguishable pre-images.
- 1134 – Structural dependencies introduced by $px(\cdot)$ further impede the isolation of
 1135 marked states necessary for Grover’s algorithm.

1136 • Error Amplification and Post-Processing Complexity:

- 1137 – Outputs of $px(\cdot)$ exhibit exponential entropy, requiring $O(2^{297})$ operations to
 1138 correlate cosets with subgroup elements.
- 1139 – Modular reductions and non-linear transformations propagate noise in quan-
 1140 tum superpositions, amplifying errors and degrading adversarial coherence.

1141 • Theoretical Structural Attacks:

- 1142 – While this form of attack is hypothetical, we anticipate a variety of advances
 1143 in topological quantum computing, enabling the next generation of quantum
 1144 algorithms based on braids, toroids, hypercubes, and other topological struc-
 1145 tures to exploit periodicity in ways we have yet to consider.
- 1146 – The uniformly unstructured nature of this construction, combined with con-
 1147 stant randomization of half of the matrix group makes this form of attack less
 1148 likely to succeed in the future.

1149 By obfuscating structural relationships and enforcing exponential search complexity,
 1150 $px(\cdot)$ ensures that recovering the hidden subgroup N remains computationally infeasible
 1151 under both classical and quantum adversarial models. These properties align quantum
 1152 complexity with classical bounds, establishing $px(\cdot)$ as a robust cryptographic primitive.

8 Information-Theoretic Security of $px(t')$

Theorem. Let $px : G \rightarrow Y$ be a mapping function with equivalence classes of size $|px^{-1}(y)| \geq 2^k$ for all $y \in Y$. Then:

1. The adversary's mutual information $I(t'; px(t'))$ is negligible, bounded by ε , where ε is a function of the compression ratio $|G|/|Y|$.
2. The adversary's probability of recovering t' from $px(t')$ is negligible, bounded by $\frac{1}{|px^{-1}(px(t'))|}$.

Proof.

Definitions and Setup. Let t' represent the hidden group elements, and $Y = px(t')$ the observed outputs. The mapping px compresses G into Y , such that each $y \in Y$ corresponds to an equivalence class of size $|px^{-1}(y)|$.

Mutual Information Bound. Mutual information is defined as:

$$I(t'; Y) = H(t') - H(t' | Y),$$

where:

- $H(t') = \log_2(|G|)$, the entropy of t' ,
- $H(t' | Y) = \log_2(|px^{-1}(y)|)$, the conditional entropy of t' given Y .

Substituting, we have:

$$I(t'; Y) = \log_2(|G|) - \log_2(|px^{-1}(y)|).$$

Rewriting in terms of the compression ratio $|G|/|Y|$, the leakage is:

$$I(t'; Y) \leq \log_2 \left(\frac{|G|}{|Y|} \right).$$

To ensure negligible leakage, the compression ratio $|G|/|Y|$ must satisfy:

$$\log_2 \left(\frac{|G|}{|Y|} \right) \leq \varepsilon,$$

where ε is a negligible function of the security parameter n .

Adversarial Success Probability. The adversary's probability of recovering t' given Y is:

$$\Pr[\text{Recover } t'] = \frac{1}{|px^{-1}(px(t'))|}.$$

Since $|px^{-1}(px(t'))| \geq 2^k$, this probability is:

$$\Pr[\text{Recover } t'] \leq 2^{-k}.$$

For sufficiently large k , this probability is negligible:

$$\Pr[\text{Recover } t'] \leq \text{negl}(n).$$

Subgroup Recovery. Recovering the hidden subgroup N requires solving the Non-Abelian Hidden Subgroup Problem (NAHSP). The adversary cannot distinguish elements in $px(t')$ without solving NAHSP, ensuring that N remains hidden.

Conclusion.

- 1180 1. The mutual information $I(t'; px(t'))$ is bounded by $\log_2(|G|/|Y|)$, which can be
 1181 made negligible by choosing sufficiently large parameters $|G|$ and $|px^{-1}(y)|$.
- 1182 2. The adversary's probability of recovering t' or N is negligible, ensuring that the
 1183 system achieves practical information-theoretic security.

1184 □

1185 9 Proof of Consistency as Verification Under Homo- 1186 morphic Transformations

1187 **Lemma 8.** *The verification equation $LHS' = RHS'$ holds if and only if the signature σ*
 1188 *is generated using the corresponding private keys and the specified public key $pk||fs$, with*
 1189 *high probability.*

1190 *Proof.* Let the key generation, signing, and verification functions be defined as follows:

1191 1. Key Generation

$$t = A \cdot x \pmod{q}, \quad t' = U \cdot t \pmod{q}, \quad pk = px(t') \pmod{q},$$

1192 where:

- 1193 • A is a public matrix,
- 1194 • x is the secret key,
- 1195 • U is a private matrix,
- 1196 • px is the mapping function.

1197 2. Signature Generation

$$\sigma = px(U \cdot (J(C1) \circ t)) \pmod{q},$$

1198 where:

- 1199 • J is a hash function (e.g., SHAKE),
- 1200 • $C1$ is constraint1, derived from $pk \cdot B$ intermediates after cubing and hashing.
- 1201 • \circ denotes a Hadamard product.

1202 3. Verification Function

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

1203 **Step 1: Valid Signature Consistency**

1204 - Substitute the signature generation equation into RHS:

$$\text{RHS} = A \cdot (px(U \cdot (J(C1) \circ t)) \circ J(C2)) \pmod q.$$

1205 - Using the properties of px and $t' = U \cdot t$, it follows that:

$$px(U \cdot (J(C1) \circ t)) = px(t') \pmod q,$$

1206 where t' satisfies the public key equation $pk = px(t') \pmod q$. - Therefore, the transfor-
1207 mations applied during signing and verification align, yielding:

$$\text{RHS}' = px(\text{RHS}) = pk \pmod q.$$

1208 **Step 2: Validating LHS'**

1209 - Substitute pk into LHS:

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

1210 - Apply the transformation px :

$$\text{LHS}' = px(\text{LHS} \circ J(C1)) \pmod q.$$

1211 - Since the signature σ was generated using the correct private key, the transformations
1212 $J(C2)$ compensate for modular inconsistencies, ensuring:

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

1213 **Step 3: Equivalence of LHS' and RHS'**

1214 - Both LHS' and RHS' reduce to $pk \pmod q$, implying:

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

1215 **Step 4: Probabilistic Argument for Invalid Signature**

1216 - For an invalid σ , the transformations in LHS and RHS will not align. To quantify this:
1217 - The output of $J(pk||fs)$ is uniformly distributed over its range. - Each σ candidate
1218 not generated with the correct private key maps to a random equivalence class under px ,
1219 with negligible probability of aligning with LHS. - The adversary must guess both:

- 1220 • σ , which depends on the secret key x and the private matrix U ,
1221 • The hash $J(pk||fs)$, which is computationally infeasible due to the pre-image resis-
1222 tance of J .

1223 - The success probability of forging σ without knowledge of x is bounded by:

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

1224 where q^n is the size of the search space for σ . This represents an information-theoretic
1225 lower bound on the success probability.

1226 **Step 5: Contrapositive**

1227 - For an invalid σ , the mismatch between LHS' and RHS' occurs due to inconsistencies
1228 in equivalence class mapping, leading to:

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

1229 **Conclusion**

1230 The verification equation $\text{LHS}' = \text{RHS}'$ holds if and only if the signature σ is generated
1231 using the valid private key x , the private matrix U , and the specified public key and basis
1232 B constraint $C1$. The probabilistic argument establishes that forging a valid σ without
1233 knowledge of the private key is computationally infeasible with high probability. \square

1234 **10 Implementation Details**

1235 **10.1 Matrix Generation Using Diverse Cryptographically Secure PRNGs**
1236

1237 To ensure cryptographic security and reproducibility, the public and private matrices in
1238 our construction should be generated deterministically using distinct cryptographically
1239 secure pseudorandom number generators (CSPRNGs). These are recommendations for
1240 high security, and certain implementations may prefer alternate functions.

1241 **10.1.1 Public Matrix Generation**

1242 The public matrices A used to generate the subgroup H_{right} are derived using AES-
1243 DRBG, per NIST-approved DRBG specifications. Each matrix $A \in \mathbb{Z}_q^{n \times n}$ is constructed
1244 as follows:

- 1245 1. Input: A 256-bit public seed Seed_A , which may be application-specific or predefined.
- 1246 2. Generation: Use AES-DRBG in CTR mode to generate n^2 entries.
- 1247 3. Mapping: Map each entry modulo q to produce a dense, full-rank matrix A .
- 1248 4. Validation: Optionally verify A 's rank to ensure it is full rank.

1249 This deterministic process is efficient, ensures reproducibility, and eliminates reliance on
1250 weak randomness.

1251 **10.1.2 Private Matrix Generation**

1252 The private matrices U , which define the subgroup H_{left} , are generated using SHA-512,
1253 SHA3-512, or SHAKE-256:

- 1254 1. Pre-Input: Optionally use a private 256-bit (or larger) value to key the hash func-
1255 tion.
- 1256 2. Input: A 256-bit private seed Seed_U , derived from an entropy source or securely
1257 exchanged during key generation.

- 1258 3. Hashing: Apply the chosen hash function to Seed_U to produce n^2 pseudorandom
1259 outputs.
- 1260 4. Mapping: Map these outputs modulo q to construct U , ensuring full rank and
1261 density.
- 1262 5. Validation: Optionally verify U 's rank to confirm full rank.

1263 10.1.3 Security Implications

1264 Using AES-DRBG for public matrices and SHA-512/SHA3/SHAKE for private matrices
1265 ensures high entropy, cryptographic security, compliance with NIST standards, and
1266 diversity in matrix generation. These methods eliminate correlations between A and U ,
1267 ensuring the subgroup structures H_{right} and H_{left} align with the theoretical reductions to
1268 NAHSP. Deterministic generation guarantees that the matrices are free from vulnerabil-
1269 ities introduced by weak or biased randomness. Furthermore, ensuring full rank for both
1270 matrices preserves the cryptographic strength of the construction.

1271 10.2 Algorithm Details

1272 10.3 Utility Algorithms

1273 Note that we don't specify the specific pseudo-random algorithm used to expand the
1274 *seed* value, as this function is designed to be modular. In our reference instance we use
1275 AES256-DRBG, but other PRNG constructions are certainly supported.

Algorithm 1 $\text{Sample}(\text{seed})$

Generates a matrix for the with non-zero elements.

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

```

1: Initialize  $A[K][N][N]$  as an empty matrix.
2: for  $\text{mat} = 0$  to  $K - 1$  do
3:    $\text{rows\_written} \leftarrow 0$ .
4:   while  $\text{rows\_written} < N$  do
5:     Generate pseudo-random buffer  $\text{buff}$  using  $\text{seed}$ .
6:     for  $y = 0$  to  $N - 1$  do
7:       Extract  $\text{trial\_vec}$  from  $\text{buff}$ .
8:       Apply transformation  $\text{NTT}(\text{trial\_vec}, Q1, R1)$ .
9:       if  $\text{trial\_vec}$  contains no zero elements then
10:        Store  $\text{trial\_vec}$  in  $A[\text{mat}][\text{rows\_written}]$ .
11:         $\text{rows\_written} \leftarrow \text{rows\_written} + 1$ .
12:        if  $\text{rows\_written} = N$  then
13:          break inner loop.
14:        end if
15:      end if
16:    end for
17:  end while
18: end for
19: return  $A$ .
```

Algorithm 2 $\text{genC2}(\text{elm1}, \text{elm2}, m, \text{SIG_fs}, \text{SIG_r})$

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

Require: Elements $\text{elm1}, \text{elm2}$ of size N , public variables $m, \text{SIG_fs}, \text{SIG_r}$ of size SEED_SIZE , and prime modulus $Q1$.

Ensure: Element $v[N]$ with non-zero elements.

- 1: Initialize SHAKE256 context: mdctx .
- 2: **if** mdctx initialization fails **then**
- 3: Throw error and terminate.
- 4: **end if**
- 5: Begin SHAKE256 hashing process.
- 6: Update hash with $\text{elm1}, \text{elm2}, m, \text{SIG_fs}$, and SIG_r .
- 7: Finalize hash to produce $\text{hash_output}[N \times \text{sizeof}(\text{int32_t})]$.
- 8: **for** $i = 0$ to $N - 1$ **do**
- 9: Extract val from $\text{hash_output}[i]$.
- 10: Compute $v[i] \leftarrow \text{abs}(\text{val}) \bmod Q1$.
- 11: **if** $v[i] = 0$ **then**
- 12: Set $v[i] \leftarrow 1$ to ensure non-zero component.
- 13: **end if**
- 14: **end for**
- 15: Free SHAKE256 context: mdctx .
- 16: **return** v .

Algorithm 3 $\text{genC1}(pk, \text{SIG_MATRIX})$

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

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

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

```
1: Initialize vector  $LHS \leftarrow pk$ .
2: Initialize  $result[N] \leftarrow 0$ ,  $LHS[N] \leftarrow pk$ .
3: for  $mat = 0$  to  $K - 1$  do
4:   Reset  $result[N] \leftarrow 0$ .
5:    $result \leftarrow \text{MatrixVectorProduct}(\text{SIG\_MATRIX}[mat], LHS, Q1)$ .
6:    $LHS \leftarrow result$ .
7:   Compute  $LHS \leftarrow LHS^3 \pmod{Q1}$ .
8:   Initialize SHAKE256 context:  $mdctx$ .
9:   if  $mdctx$  initialization fails then
10:    Throw error and terminate.
11:  end if
12:  Begin SHAKE256 hashing process.
13:  Update hash with LHS, result.
14:  Finalize hash to produce  $hash\_output[N]$ .
15:  for  $i = 0$  to  $N - 1$  do
16:    Extract  $val$  from  $hash\_output[i]$ .
17:    Compute  $v[i] \leftarrow \text{abs}(val) \pmod{Q1}$ .
18:    if  $v[i] = 0$  then
19:      Set  $v[i] \leftarrow 1$  to ensure non-zero component.
20:    end if
21:  end for
22:  Free SHAKE256 context:  $mdctx$ .
23:  Store  $LHS$  in  $C1[mat]$ .
24: end for
25: return  $C1$ .
```

Table 4: 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.

- 1: Initialize secretKeys[K][N] uniformly random x in the range $[1, 255]$.
 - 2: Initialize matrix: MATRIX_A[K][N][N].
 - 3: Initialize matrix: MATRIX_U[K][N][N].
 - 4: Generate random seed: PK_SEED_A of length SEED_SIZE.
 - 5: Generate random seed: SK_SEED_U of length SEED_SIZE.
 - 6: Sample MATRIX_A using PK_SEED_A.
 - 7: Sample MATRIX_U using SK_SEED_U.
 - 8: Initialize current_pk \leftarrow secretKeys[0].
 - 9: current_pk \leftarrow NTT(current_pk, $Q1$, $R1$).
 - 10: **for** $I = 0$ to $K - 1$ **do** \triangleright Iterate through the chain of transformations.
 - 11: Initialize result[N] \leftarrow 0.
 - 12: **if** $I > 0$ **then**
 - 13: Update skey \leftarrow secretKeys[I] and compute skey \leftarrow NTT(skey, $Q1$, $R1$).
 - 14: Element-wise multiplication: current_pk \leftarrow skey \circ current_pk mod $Q1$.
 - 15: **end if**
 - 16: result \leftarrow MatrixVectorProduct(MATRIX_A[I], current_pk, $Q1$).
 - 17: current_pk \leftarrow result.
 - 18: result[N] \leftarrow 0.
 - 19: result \leftarrow MatrixVectorProduct(MATRIX_U[I], current_pk, $Q1$).
 - 20: Apply hiding function: current_pk \leftarrow px (result).
 - 21: **end for**
 - 22: Ensure non-zero condition: nonzero_count(current_pk) $\geq N$.
 - 23: **if** Condition fails **then**
 - 24: Retry key generation.
 - 25: **end if**
 - 26: **return** current_pk, PK_SEED_A, SK_SEED_U, secretKeys[K]
-

10.4 Signature Generation

Algorithm 5 $\text{Sign}(m, \text{secretKeys}[K], \text{PK_SEED_A}, \text{pk_elem}, \text{SK_SEED_U})$

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 $\text{secretKeys}[K][N]$, public seed PK_SEED_A , public key pk_elem , secret seed SK_SEED_U .

- 1: Initialize $\text{sig}[N] \leftarrow 0$.
- 2: Set $\text{SIG_COMPLETED} \leftarrow 0$.
- 3: Generate random bytes: $\text{rand_A}, \text{rand_B}$ of size $SEED_SIZE$.
- 4: Compute FS using $\text{shake256}(\text{rand_A}, \text{PK_SEED_A}, \text{pk_element})$.
- 5: Compute SIG_SEED_B using $\text{shake256}(\text{FS}, m, \text{rand_B}, \text{pk_element})$.
- 6: Sample matrix: $\text{MATRIX_B}[K][N][N]$ using SIG_SEED_B .
- 7: Sample matrix: $\text{MATRIX_U}[K][N][N]$ using SK_SEED_U .
- 8: Initialize $\text{C1}[K][N]$ via $\text{genC1}(\text{pk}, \text{MATRIX_B})$.
- 9: Set $\text{sig} \leftarrow \text{secretKeys}[0]$ and apply $\text{forward_ntt}(\text{sig}, Q1, R1)$.
- 10: **for** $I = 0$ to $K - 1$ **do**
- 11: **if** $I > 0$ **then**
- 12: $\text{skey} \leftarrow \text{forward_ntt}(\text{secretKeys}[I], Q1, R1)$.
- 13: $\text{sig} \leftarrow \text{skey} \circ \text{sig} \pmod{Q1}$.
- 14: **end if**
- 15: $\text{result}[N] \leftarrow \text{MatrixVectorProduct}(\text{MATRIX_B}[I], \text{sig}, Q1)$.
- 16: $\text{sig} \leftarrow \text{result}$
- 17: $\text{result}[N] \leftarrow 0$.
- 18: $\text{result}[N] \leftarrow \text{MatrixVectorProduct}(\text{MATRIX_U}[I], \text{sig}, Q1)$.
- 19: $\text{sig} \leftarrow \text{sig} \circ \text{C1}[I] \pmod{Q1}$.
- 20: Apply Hiding Function $\text{sig} \leftarrow \text{px}(\text{result})$.
- 21: **end for**
- 22: Apply $\text{inverse_ntt}(\text{sig}, Q1, R1)$.
- 23: Validate Entropy $\text{C3_CHECK} \leftarrow \text{Verify_entropy}(\text{sig})$.
- 24: C3 Check Retry after clearing buffers if $\text{C3_CHECK} = 1$.
- 25: Count non-zero elements in sig .
- 26: **if** $\text{nonzero_count}(\text{sig}) \geq N$ **then**
- 27: Output $\text{sig}, \text{FS}, \text{rand_B}$.
- 28: **else**
- 29: Retry after clearing buffers.
- 30: **end if**

10.5 Signature Verification

1279 This function uses the Fiat-Shamir heuristic to reconstruct the signature basis seed that
 1280 was used during the signing process. Together with the public key matrix seed, both
 1281 public and signature matrices are sampled and 'swapped', such that sig signature element
 1282 is transformed by the public basis and the public key element is transformed by the
 1283 signature basis. The public key is also isomorphically transformed by the masking value
 1284 that was used at each layer during signing. This mask is derived from the interaction
 1285 between the public key and the signature basis, effectively binding them together.

1286 Additionally as valid signatures are information theoretically guaranteed to have ob-
 1287 servational 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 $rand_B$.

- 1: Initialize $SIG_SEED_B[SEED_SIZE]$, $C2[N]$, $C3_CHECK \leftarrow 0$.
 - 2: Validate Entropy $C3_CHECK \leftarrow Verify_entropy(sig)$.
 - 3: C3 Check **return** 0 if $C3_CHECK = 1$.
 - 4: Compute $C2$ using $genC2(sig, pk_elem, m, FS, rand_B)$.
 - 5: Apply $forward_ntt(C2, Q1, R1)$ and $forward_ntt(sig, Q1, R1)$.
 - 6: Construct $temp_fs$ by concatenating FS , m , $rand_B$, and pk_elem .
 - 7: Compute SIG_SEED_B using $shake256(temp_fs)$.
 - 8: Sample matrix: $PK_MATRIX[K][N][N]$ using PK_SEED_A
 - 9: Sample matrix: $SIG_MATRIX[K][N][N]$ using SIG_SEED_B .
 - 10: Initialize $LHS[N] \leftarrow pk_elem$, $RHS[N] \leftarrow sig$.
 - 11: Compute $C1[K][N]$ using $genC1(pk, SIG_MATRIX)$.
 - 12: **for** $I = 0$ to $K - 1$ **do**
 - 13: $LHS \leftarrow LHS \circ C2 \pmod{Q1}$.
 - 14: $LHS \leftarrow MatrixVectorProduct(SIG_MATRIX[I], LHS)$.
 - 15: $LHS \leftarrow LHS \circ C1[I] \pmod{Q1}$.
 - 16: **if** $I \neq K - 1$ **then**
 - 17: Apply Equivocation function: $LHS \leftarrow px(LHS)$.
 - 18: **end if**
 - 19: **end for**
 - 20: **for** $I = 0$ to $K - 1$ **do**
 - 21: $RHS \leftarrow RHS \circ C2 \pmod{Q1}$.
 - 22: $RHS \leftarrow MatrixVectorProduct(PK_MATRIX[I], RHS)$.
 - 23: **if** $I \neq K - 1$ **then**
 - 24: Apply Equivocation function: $RHS \leftarrow px(RHS)$.
 - 25: **end if**
 - 26: **end for**
 - 27: Compare RHS and LHS .
 - 28: **return** 0 if equal, otherwise 1.
-

1288 **Observed Entropy Rejection Sampling**

1289 As part of what we are calling the third constraint, we want to rejection sample based on
 1290 the observed randomness of the σ input. We are considering elements of length of $n =$
 1291 $\{64, 128, 256\}$ and have implemented constraints thus far on two probabilistic features,
 1292 bit level and byte level randomness. In the byte probability case, we are measuring each
 1293 component byte value, with the ideal σ having zero colliding component values. However,
 1294 we find that given the small sets of we are considering, we see multiple values appear two
 1295 or three times, despite being valid.

1296 To increase the granularity of randomness checking, we measure the raw ratio of 0 and
 1297 1 value bits across the array. For $n = 128$, where an ideally random signature would have

1298 512 value 0 bits and 512 value 1 bits. To constrain input signatures to approximately
 1299 1/10 of the total possible modular space, we set the threshold ratio to 0.991. But,
 1300 a "downshifted" element where components were in the range of $\{0, \dots, 255\}$ entirely
 1301 consisting of value 128 would pass with 512 0 bits and 512 1 bits. To correctly reject
 1302 invalid signatures, we measure both bit level entropy and byte level entropy.

Table 5: Observed Entropy Thresholds for Different Values of using Byte probabilities N

N	H_THRESHOLD
64	5.8
128	6.8
256	7.8

Table 6: 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 .

- 1: Extract entropy-relevant components from sig : $SIG_VALUES \leftarrow \text{extract}(sig)$.
- 2: Compute the empirical distribution $DIST$ of SIG_VALUES over the modular domain $[0, Q)$.
- 3: Calculate the Shannon entropy H_SIG :

$$H_SIG \leftarrow - \sum_{x \in DIST} p(x) \log p(x),$$

where $p(x)$ is the probability of x in $DIST$.

- 4: Extract count of 0 bits and 1 bits from sig as $COUNT0$ and $COUNT1$.
 - 5: Compute Ratio as HB_SIG .
 - 6: **if** $H_SIG < H_THRESHOLD$ and $HB_SIG < HB_THRESHOLD$ **then**
 - 7: **return** 1 ▷ Entropy too low, validation fails.
 - 8: **else**
 - 9: **return** 0 ▷ Entropy validation succeeds.
 - 10: **end if**
-

1303 While the underlying concept of rejection sampling based on entropy should be clear,
 1304 the exact implementation is to be refined in subsequent revisions of this preprint. The
 1305 achievable constraint is that the adversary cannot forge σ using the entire ambient mod-
 1306 ular space, but only a specific fraction of it. This will lead to more refined and accurate
 1307 probability.

10.6 Hiding Function

Table 7: Hiding Function NTT Parameters - Prime Fields and ω Roots of Unity

Level	Q1	R1	Q2	R2	R3
I	257	81	1283	3	3
III	257	9	1283	3	3
V	257	3	1283	3	3

Each round of modular addition causes approximately half of the coefficients to wrap around the modulus, creating information loss and diffusion. Mapping to a higher modulus, combined with scaling and permutation is responsible for the bulk of pre-images introduced by the hiding function. This function can be modified based on a predetermined optimal mapping ratio of output to potential valid pre-image inputs. As an initial setting, for level III, we target a compression ratio of 2^{293} to 1 out of q^n possibilities, creating a computationally sufficient number of indistinguishable elements per coset, invariant across both the hidden and ambient group elements.

Algorithm 8 Hiding Function $px(\text{in_elem})$

Transforms an input, a coset of group N , altering the structure using a series of modular operations creating a many-to-one mapping.

Require: Input vector $\text{in_elem}[N]$, prime moduli $Q1, Q2$, roots $R1, R2$.

Ensure: Output vector $\text{out_elem}[N]$.

- 1: Allocate temporary arrays: $\text{vecsq}[N], \text{vec}[n]$.
 - 2: Copy $\text{vecsq} \leftarrow \text{in_elem}$.
 - 3: $\text{vecsq} \leftarrow \text{inverse_ntt}(\text{vecsq}, Q1, R1)$.
 - 4: $\text{vecsq} \leftarrow \text{forward_ntt}(\text{vecsq}, Q2, R2)$.
 - 5: Copy $\text{vec} \leftarrow \text{vecsq}$.
 - 6: $\text{vecsq} \leftarrow \text{pointwise_addition}(\text{vecsq}, \text{vecsq}, Q2)$.
 - 7: $\text{vecsq} \leftarrow \text{pointwise_addition}(\text{vecsq}, \text{vec}, Q2)$.
 - 8: $\text{vecsq} \leftarrow \text{pointwise_addition}(\text{vecsq}, \text{vec}, Q2)$.
 - 9: $\text{vecsq} \leftarrow \text{pointwise_addition}(\text{vecsq}, \text{vec}, Q2)$.
 - 10: $\text{vecsq} \leftarrow \text{inverse_ntt}(\text{vecsq}, Q2, R2)$.
 - 11: $\text{vecsq} \leftarrow \text{forward_ntt}(\text{vecsq}, Q1, R1)$.
 - 12: $\text{out_vec} \leftarrow \text{vecsq}$
 - 13: **return** out_vec .
-

Note that step 11 performs an NTT forward transformation in the $Q1$ domain. After step 10, as a result of the internal diffusion and inverse NTT from $Q2$, we will have an element, (for level III) will look like this:

320, 925, 1060, 280, 475, 730, 1065, 20, 1120, 845, 1110, 125, 255, 430, 1245, 230, 165, 300, 695, 565,
980, 1005, 175, 860, 1135, 785, 610, 765, 760, 855, 1000, 175, 485, 310, 635, 710, 1015, 1110, 240, 1155, 40,
960, 1225, 840, 545, 220, 1005, 390, 940, 765, 1245, 40, 840, 320, 750, 1170, 120, 410, 480, 1270, 530, 470,
380, 530, 290, 25, 350, 325, 1265, 1005, 1275, 0, 975, 1055, 315, 1005, 915, 985, 240, 545, 455, 730, 570,
875, 320, 60, 200, 835, 880, 1205, 685, 1190, 1200, 495, 260, 245, 300, 370, 120, 700, 795, 330, 295, 705,

1324 660, 695, 320, 455, 905, 1095, 105, 300, 30, 145, 1095, 900, 285, 1010, 395, 650, 695, 465, 1195, 545, 1185,
1325 305, 255, 1175, 128

1326 This is a result of using an NTT ω that isn't technically 'valid' for this field and array
1327 size. This is fine and intentional. We aren't performing any convolutional operations
1328 with this field (so the ω value isn't a factor) and it gives us a uniformly distributed set
1329 of components across 0 to 1282. We perform step 11 to explicitly alias the array back to
1330 the same residue class as the $q = 257$ NTT domain. After step 11, for level III, the same
1331 array becomes:

253, 85, 34, 171, 146, 12, 128, 73, 193, 243, 20, 119, 243, 158, 19, 237, 198, 185, 149, 15, 0, 0, 88, 112, 0,
1332 193, 0, 201, 181, 68, 131, 0, 216, 14, 96, 130, 206, 220, 168, 153, 251, 200, 194, 154, 116, 151, 59, 163, 152,
1333 193, 228, 198, 107, 170, 243, 111, 117, 209, 215, 230, 145, 165, 29, 252, 200, 82, 52, 115, 241, 55, 221,
1334 135, 137, 243, 119, 197, 169, 216, 162, 70, 196, 234, 236, 178, 201, 66, 114, 101, 249, 26, 4, 184, 30, 76,
1335 218, 201, 97, 240, 186, 182, 83, 248, 184, 7, 162, 178, 90, 94, 205, 43, 69, 213, 234, 55, 131, 253, 49,
1336 208, 167, 216, 130, 108, 18, 142, 155, 205, 137, 217, 128

1337 This aliasing of components from a field ≈ 5 times larger down to $q = 257$ is how we
1338 generate multiple pre-images per output, and compress inputs to a specific equivalence
1339 class.

1340 Empirical testing demonstrates that outputs from the mapping function $px(\cdot)$ ap-
1341 pear statistically uniform and indistinguishable from random values. This uniformity is
1342 achieved through the interplay of NTT-based projections, modular scaling, and iterative
1343 mixing operations applied over a larger finite field.

1344 The vector t spans the full ambient group G as matrix A is full rank and changes with
1345 each operation and secret x is uniformly random. combined This projection $t' \equiv U \cdot t$ also
1346 spans the entire ambient space. This process effectively blurs the cosets of N , distributing
1347 them uniformly over Z_n^q and mapped to a specific equivalence classes created by $px(\cdot)$.

1348 As a result, recovering all valid indistinguishable pre-images for a given output element
1349 of $px(\cdot)$ is insufficient to reconstruct N . The pre-images include a superset of elements
1350 comprising valid members of unrelated ambient group elements, with probabilistically
1351 only one valid element. This amalgamation obscures the boundaries of the 'true' N ,
1352 making it computationally infeasible to distinguish subgroup membership based solely
1353 on inversion of $px(\cdot)$.

1354 11 Attack Models

1355 In this section, we rigorously analyze the security of the proposed cryptographic scheme
1356 against both classical and quantum adversaries. We focus on proving that the scheme
1357 achieves IND-CPA (Indistinguishability under Chosen Plaintext Attack) security by demon-
1358 strating the computational infeasibility of recovering the hidden subgroup N or distin-
1359 guishing ciphertexts under the specified attack models.

11.1 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.

11.1.1 Preliminaries

Let us recall the key components:

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

11.1.2 Proof of Security Against Classical Adversaries

Lemma 1 (Computational Indistinguishability). *Under the assumption that $px(\cdot)$ 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 $f(g) = px(U \cdot A \cdot x)$.
- As established in Section 3, solving NAHSP in this group is computationally infeasible for classical adversaries.

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

- The function $px(\cdot)$ introduces high ambiguity, mapping exponentially many inputs to the same output.
- From Lemma 3 in Section 4.3, we know that the outputs of $px(\cdot)$ 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 $px(U \cdot A \cdot x)$ 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.

11.2 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.

11.2.1 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 $px(\cdot)$, 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 $G = H_{\text{left}} \times H_{\text{right}}$ 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 $px(\cdot)$.

- The mapping function $px(\cdot)$ further complicates any attempt to extract information via quantum algorithms.

- From Lemma 7 in Section 4.7, even if a quantum adversary could invert $px(\cdot)$, 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 $px(\cdot)$, quantum adversaries cannot break the scheme with non-negligible probability. Therefore, the scheme achieves IND-CPA security even in the presence of quantum adversaries.

11.3 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 $px(\cdot)$ behaves as a pseudorandom function, the proposed digital signature scheme is IND-CPA secure.*

1438 **Proof. Definition of IND-CPA Security.**

1439 A digital signature scheme is IND-CPA secure if no polynomial-time adversary can
1440 distinguish between signatures of chosen messages, even when given access to a signing
1441 oracle.

1442 **Game-Based Proof Structure.**

1443 We consider the standard IND-CPA security game between a challenger and an ad-
1444 versary \mathcal{A} :

1445 1. **Setup:** The challenger generates a public-private key pair (pk, sk) and provides
1446 pk to \mathcal{A} .

1447 2. **Query Phase:** \mathcal{A} may request signatures on messages of its choice.

1448 3. **Challenge Phase:** \mathcal{A} selects two messages m_0 and m_1 . The challenger randomly
1449 selects $b \in \{0, 1\}$ and returns $\sigma = \text{Sign}(m_b, sk)$.

1450 4. **Guess Phase:** \mathcal{A} outputs a guess b' . The adversary wins if $b' = b$.

1451 Our goal is to show that $\Pr[b' = b] \leq \frac{1}{2} + \varepsilon$, where ε is negligible.

1452 **Analysis.**

1453 Assume, for contradiction, that \mathcal{A} can win the game with a non-negligible advantage
1454 δ .

1455 **Step 1: Construction of a Simulator to Solve NAHSP.**

1456 We construct a simulator \mathcal{S} that uses \mathcal{A} to solve the NAHSP:

1457 - \mathcal{S} receives an instance of NAHSP in G and needs to find the hidden subgroup N .

1458 - \mathcal{S} simulates the challenger for \mathcal{A} , using the NAHSP instance to generate public keys
1459 and signatures.

1460 - When \mathcal{A} outputs b' , \mathcal{S} uses this information to extract information about N .

1461 **Step 2: Contradiction with the Hardness of NAHSP.**

1462 - If \mathcal{S} can solve NAHSP using \mathcal{A} 's advantage δ , then the hardness assumption of
1463 NAHSP is violated.

1464 - Therefore, δ must be negligible.

1465 **Step 3: Security Reduction via Hybrid Arguments.**

1466 - We can define a sequence of hybrid experiments transitioning from the real scheme
1467 to an ideal scheme where signatures are replaced with random values.

1468 - The indistinguishability of outputs from $px(\cdot)$ ensures that \mathcal{A} cannot distinguish
1469 between hybrids with non-negligible advantage.

1470 **Conclusion.**

1471 Since any non-negligible advantage δ leads to a contradiction with the hardness of
1472 NAHSP, we conclude that \mathcal{A} cannot win the IND-CPA game with more than negligible
1473 advantage. Therefore, the scheme is IND-CPA secure.

1474 **11.4 Resistance to Forgery Under Chosen Message Attacks**

1475 **Theorem 2 (Unforgeability under Chosen Message Attack).** *Assuming the hard-*
1476 *ness of NAHSP and the collision resistance of the hash function $J(\cdot)$, the proposed scheme*
1477 *is existentially unforgeable under chosen message attacks (EUF-CMA).*

1478 **Proof. Step 1: Assumptions and Definitions.** - Let \mathcal{A} be an adversary attempting
1479 to forge a valid signature σ^* for a message m^* that has not been queried during the
1480 signing oracle phase. - The scheme uses the Fiat-Shamir heuristic to bind the signature
1481 to the message, the public key, and a random nonce. - A valid forgery requires \mathcal{A} to

1482 produce σ^* such that:

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

1483 **Step 2: Connection to NAHSP and $px(\cdot)$.** - To forge σ^* , \mathcal{A} must either: 1.
1484 Recover the private key x or the hidden matrix U , allowing the computation of valid
1485 transformations. This is equivalent to solving the Non-Abelian Hidden Subgroup Problem
1486 (NAHSP), which is assumed to be hard. 2. Generate a valid preimage under $px(\cdot)$
1487 without access to the private key or matrix. The lossy, many-to-one nature of $px(\cdot)$
1488 ensures that the adversary cannot distinguish valid preimages from an exponentially
1489 large indistinguishable set.

1490 **Step 3: Resistance to Hash Function Collisions.** - The Fiat-Shamir heuristic
1491 involves the hash function $J(\cdot)$, which produces a binding challenge for the signature.
1492 For \mathcal{A} to forge σ^* , it must either: 1. Find a collision $J(pk \parallel m^* \parallel r) = J(pk \parallel m' \parallel r')$,
1493 which is infeasible due to the assumed collision resistance of $J(\cdot)$. 2. Guess the challenge
1494 generated by $J(\cdot)$ and align it with a valid subgroup element. The probability of such a
1495 guess is negligible due to the high entropy of the output space of $J(\cdot)$.

1496 **Step 4: Reduction to a Hard Problem.** - Assume \mathcal{A} successfully forges σ^* with
1497 non-negligible probability. We construct a simulator \mathcal{S} that uses \mathcal{A} to solve NAHSP or
1498 find a collision in $J(\cdot)$: 1. \mathcal{S} simulates the signing oracle for \mathcal{A} , generating signatures
1499 using a secret key x and the private matrix U . 2. If \mathcal{A} outputs a valid forgery σ^* , \mathcal{S}
1500 uses σ^* to extract information about the hidden subgroup N or to find a collision in $J(\cdot)$.
1501 3. Since both outcomes contradict the hardness of NAHSP or the collision resistance of
1502 $J(\cdot)$, \mathcal{A} 's success probability must be negligible.

1503 **Step 5: Conclusion.** - The adversary \mathcal{A} cannot forge a valid signature σ^* on a
1504 message m^* without solving NAHSP, inverting $px(\cdot)$, or finding a collision in $J(\cdot)$, all of
1505 which are computationally infeasible. - Therefore, the proposed scheme is existentially
1506 unforgeable under chosen message attacks (EUF-CMA).

1507 11.5 Inapplicability of CCA2 Security

1508 It is important to note that our proposed digital signature scheme does not incorporate
1509 a decryption oracle, as it is not designed to handle encrypted messages or ciphertext
1510 directly. The absence of such an oracle renders the chosen ciphertext attack (CCA2)
1511 model irrelevant for this construction.

1512 Instead, the security of the scheme is analyzed under the IND-CPA (Indistinguishability
1513 under Chosen Plaintext Attack) and EUF-CMA (Existential Unforgeability under
1514 Chosen Message Attack) models, which are sufficient and appropriate given the nature
1515 of the signature application.

1516 By excluding a decryption oracle, the scheme eliminates a common attack vector
1517 associated with adaptive adversaries in CCA2 scenarios, further solidifying its robustness
1518 in practical cryptographic deployments.

1519 11.5.1 Brute Force Key Recovery

1520 Brute force recovery of secrets is implementation dependent. The scheme covered in
1521 this document leverages a single private matrix seed for each instance in the chain, and
1522 unique password x elements per instance. Based on size and security concerns, if private
1523 key size was critical, the private key, in theory, could be shrunk to a single secret
1524 seed that expanded to provide every hidden matrix and x element. If absolute security

1525 were paramount, each hidden matrix could be derived from its own seed, or stored fully
 1526 instantiated. In this brief analysis, we will simply derive the costs of brute forcing the
 1527 scheme as described. Each level has k chains, with one x secret element, effectively n
 1528 bytes long. Additionally, each level has one hidden matrix seed, of $SEED_SIZE$ length,
 1529 in bytes. Relative sizes are listed below:

Table 8: 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}

1530 As even level I requires 2^{4224} classical operations, brute force attacks do not appear to
 1531 be a practical concern with the variant as described in this paper. Note that this could
 1532 change depending on various implementation optimizations.

1533 11.6 Conclusion on Attack Models

1534 Through rigorous proofs, we have established that the proposed cryptographic scheme is
 1535 secure against both classical and quantum adversaries. The security relies on:

- 1536 • The computational hardness of the NAHSP in non-abelian groups.
- 1537 • The pseudorandomness and computational indistinguishability introduced by the
 1538 mapping function $px(\cdot)$.
- 1539 • The collision resistance of the hash function used in the Fiat-Shamir heuristic.

1540 These properties collectively ensure that adversaries cannot feasibly recover private
 1541 keys, forge signatures, or distinguish ciphertexts, thereby achieving IND-CPA security
 1542 and resisting forgery under chosen message attacks.

1543 12 Implementation and Efficiency

1544 12.1 Performance Evaluation

Table 9: 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

1545 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 10: Public Key, Signature Sizes, and Chain Instances Across Levels

1546 12.2 Comparison with Other Schemes

Table 11: 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

1547 The above data was been gathered from the PQShield Post-Quantum signatures zoo, we
1548 have not verified it and the most recent developments may not be reflected, but we feel
1549 the source is accurate for this early comparison. To date, our priority has been theo-
1550 retical security and communication efficiency, with little attention paid to performance
1551 optimization. Over time, it is almost certain our performance numbers will improve be-
1552 yond our reference instance. We are using portable ANSI C, openssl/TLS, our portion
1553 of the code does not leverage intrinsics, is single threaded, and the majority of time is
1554 currently spent expanding matrices from seed. With dedicated effort we feel performance
1555 and optimization will yield significant improvements. That said, we do not expect to
1556 outperform ML-DSA in terms of compute, nor do we expect computational performance
1557 to be a barrier for adoption.

⁰Data source: <https://pqshield.github.io/nist-sigs-zoo/c>

12.3 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 $\{0, \dots, 256\}$, 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.

13 Open Research Questions

1. For function $px()$, 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.

14 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 communication 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.

14.1 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.

- 1595 • **Public Key Agreement:** Enables fast, quantum-resilient key exchanges with minimal
1596 communication overhead, ideal for constrained networks.
- 1597 • **Identity-Based and Attribute-Based Cryptography:** Supports fine-grained access con-
1598 trol, allowing secure communication based on user identities or attributes without
1599 requiring heavy key distribution infrastructure.
- 1600 • **Zero-Knowledge Proofs (ZKPs):** Facilitates privacy-preserving verification of state-
1601 ments without exposing underlying secrets, essential for regulatory compliance and
1602 secure interactions.

1603 **14.2 Efficient Communication and Deployability**

1604 **Compact Communication Sizes:**

- 1605 • NAHSP-based cryptography achieves extremely small communication footprints,
1606 with signatures and keys often requiring a fraction of the size used by lattice-
1607 based systems. This efficiency is critical in bandwidth-limited environments such
1608 as undersea and RF networks.
- 1609 • **Example:** For a security level equivalent to Dilithium-III, the NAHSP-based scheme
1610 offers signatures of 80 bytes and public keys under 100 bytes, compared to the
1611 hundreds or thousands of bytes required by Dilithium.

1612 **Software-Only Deployment:**

- 1613 • Unlike lattice-based systems, which often require specialized hardware for efficient
1614 operation, NAHSP-based cryptography can be implemented as a drop-in replace-
1615 ment via software patches.
- 1616 • This allows immediate deployment in existing infrastructures, including mobile de-
1617 vices, routers, and IoT systems, without the need for hardware upgrades.
- 1618 • Rapid updates ensure forward compatibility with evolving security standards while
1619 minimizing deployment costs.

1620 **14.3 Secure Communication Across Diverse Environments**

1621 NAHSP-based cryptography's compact and efficient design enables secure communication
1622 in challenging and bandwidth-constrained domains:

- 1623 • **Cellular Networks:** Ensures efficient and secure handshakes, even in low-latency 5G
1624 environments, where minimal communication overhead is critical.
- 1625 • **Radio Frequency (RF) Systems:** Compact key sizes and signatures reduce transmis-
1626 sion time in RF-constrained settings, such as military radios and satellite uplinks.
- 1627 • **Undersea Acoustics:** Low-bandwidth undersea acoustic networks benefit from NAHSP's
1628 compact communication, enabling secure exchanges where data rates are severely
1629 limited.

1630 **14.4 Comparison with Dilithium and Other Systems**

1631 A matrix based NAHSP scheme addresses several limitations of Dilithium and similar
1632 lattice-based approaches:

- 1633 • **Smaller Communication Sizes:** Signatures and keys are significantly more compact,
1634 reducing storage and bandwidth requirements.
- 1635 • **Flexibility Across Environments:** Performs robustly in environments where lattice-
1636 based schemes face challenges, such as RF and undersea communication.
- 1637 • **Advanced Constructions:** Offers natural support for identity-based encryption and
1638 zero-knowledge proofs, features that are generally not feasible to implement over
1639 schemes with noise.

1640 **14.5 Real-World Applications**

1641 **Critical Infrastructure and Defense:**

- 1642 • Secures command and control systems in military and intelligence operations, en-
1643 suring quantum resilience and adaptability across RF and satellite links.
- 1644 • Protects SCADA systems in critical infrastructure, such as energy grids and trans-
1645 portation networks, with lightweight and efficient cryptographic primitives.

1646 **IoT and Edge Devices:**

- 1647 • Provides secure authentication for IoT devices with constrained processing power,
1648 mitigating risks of botnet attacks and data breaches.
- 1649 • Ensures efficient encryption and signing for edge devices in industrial and healthcare
1650 settings.

1651 **Blockchain and Distributed Systems:**

- 1652 • Enhances consensus mechanisms with compact, quantum-resistant signatures, re-
1653 ducing energy consumption and improving scalability.
- 1654 • Secures smart contracts and cryptographic tokens with lightweight, efficient con-
1655 structions.

1656 **Telecommunications and Financial Systems:**

- 1657 • Enables secure mobile payment systems and digital banking with minimal commu-
1658 nication overhead, ensuring transaction authenticity and integrity.
- 1659 • Modernizes public key infrastructure (PKI) for quantum resilience while minimizing
1660 deployment costs.

14.6 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 $\text{px}()$ 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.

15 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 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:

- 1701 • Cryptography: The NAHSP framework invites exploration into additional construc-
1702 tions such as group-based encryption, secure multi-party computation, and ad-
1703 vanced privacy-preserving protocols.
- 1704 • Complexity Theory: By leveraging non-abelian group properties, the system pro-
1705 vides fertile ground for studying alternate hardness assumptions and their implica-
1706 tions for classical and quantum computational limits.
- 1707 • Quantum Algorithm Design: The inherent resilience of the scheme challenges re-
1708 searchers to explore novel quantum algorithms capable of addressing non-abelian
1709 group problems, advancing our understanding of quantum computational power.
- 1710 • Systems Security: With its adaptability to constrained environments such as IoT,
1711 RF, and undersea acoustics, this system sets the stage for breakthroughs in secure
1712 communication under extreme conditions.

1713 **Closing Statement:** This cryptographic scheme offers a significant contribution to
1714 the evolving landscape of post-quantum security. By providing practical information-
1715 theoretic guarantees and addressing key implementation challenges, it has the potential
1716 to transform how secure systems are designed and deployed. While further research and
1717 optimization remain, this work lays a strong foundation for future innovations in cryp-
1718 tography, complexity theory, and quantum algorithm design, positioning it as a critical
1719 component in the journey toward resilient and scalable global security systems.

1720 16 Acknowledgements

1721 This work is provisionally patented, USPTO 63/726,609.

1722 References

- 1723 [1] Daniel J. Bernstein et al. “The SPHINCS+ Signature Framework”. In: *Proceedings*
1724 *of the 2019 ACM SIGSAC Conference on Computer and Communications Secu-*
1725 *rity*. CCS ’19. London, United Kingdom: Association for Computing Machinery,
1726 2019, pp. 2129–2146. ISBN: 9781450367479. DOI: 10.1145/3319535.3363229. URL:
1727 <https://doi.org/10.1145/3319535.3363229>.
- 1728 [2] Léo Ducas et al. “Dilithium: A high-speed lattice-based digital signature scheme”.
1729 In: *CCS 2019*. 2019, pp. 897–918.
- 1730 [3] Luca De Feo et al. *SQISign: compact post-quantum signatures from quaternions*
1731 *and isogenies*. Cryptology ePrint Archive, Paper 2020/1240. 2020. URL: <https://eprint.iacr.org/2020/1240>.
1732
- 1733 [4] David Garber. “Braid group cryptography”. In: *Braids: Introductory lectures on*
1734 *braids, configurations and their applications*. World Scientific, 2010, pp. 329–403.
- 1735 [5] Dimitri Grigoriev and Ilia Ponomarenko. *Constructions in public-key cryptography*
1736 *over matrix groups*. 2005. arXiv: math/0506180 [math.GR]. URL: <https://arxiv.org/abs/math/0506180>.
1737

- 1738 [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. 1739 2000, pp. 166–183. 1740 1741
- 1742 [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>. 1743 1744
- 1745 [8] Karl Mahlbürg. “An overview of braid group cryptography”. In: *preprint* (2004).
- 1746 [9] Alexei Myasnikov, Vladimir Shpilrain, and Alexander Ushakov. “A practical at- 1747 tack on a braid group based cryptographic protocol”. In: *Advances in Cryptology— 1748 CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, 1749 California, USA, August 14–18, 2005. Proceedings 25*. Springer. 2005, pp. 86–96.
- 1750 [10] Alexei G Myasnikov, Vladimir Shpilrain, and Alexander Ushakov. *Non-commutative 1751 cryptography and complexity of group-theoretic problems*. 177. American Mathemat- 1752 ical Soc., 2011.
- 1753 [11] Michael Schmid et al. *Falcon Takes Off - A Hardware Implementation of the Falcon 1754 Signature Scheme*. Cryptology ePrint Archive, Paper 2023/1885. 2023. URL: <https://eprint.iacr.org/2023/1885>. 1755
- 1756 [12] Claude E. Shannon. “Communication theory of secrecy systems”. In: *Bell Syst. 1757 Tech. J.* 28.4 (1949), pp. 656–715. DOI: 10.1002/J.1538-7305.1949.TB00928.X. 1758 URL: <https://doi.org/10.1002/j.1538-7305.1949.tb00928.x>.
- 1759 [13] Claude Elwood Shannon. “A Mathematical Theory of Communication”. In: *The 1760 Bell System Technical Journal* 27 (1948), pp. 379–423. URL: [http://plan9.bell- 1761 labs.com/cm/ms/what/shannonday/shannon1948.pdf](http://plan9.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf) (visited on 04/22/2003).