

The Silent Radix Is Not a Key: A Security Analysis of Unknown-Radix Encoding as a Cryptographic Primitive

The Silent Radix Is Not a Key: A Security Analysis of Unknown-Radix Encoding as a Cryptographic Primitive

Abstract

Positional numerals do not internally specify their own base, and a recent informal proposal — Silent Radix Encryption (SRE) — conjectures that this "silent radix" can serve as a cryptographic secret: plaintext symbols are mapped to digits by a secret map and read as an integer in a secret radix B , with B acting as a trapdoor for key encapsulation. No formal security analysis of this construction exists. We formalize the encoding core, define its key space over admissible radices $B \in [2, 2^{16}]$ and digit maps on a 26-symbol alphabet, and compute its entropy: 104.38 bits. We then show the construction reduces to a monoalphabetic substitution cipher with base-dependent segmentation, and we derive a two-stage ciphertext-only attack: a histogram test recovers B at a cost of at most 6.5535×10^7 digit operations for a 1000-symbol ciphertext, after which standard frequency analysis applies within a unicity distance of 149.1 symbols, computed from the key entropy and an assumed language redundancy of 0.70 bits/symbol. Brute force over the full key space at 10^9 keys/s would take 8.36×10^{14} years, yet the structural attack succeeds on a single short message: key-space size is irrelevant. We conclude that radix obscurity offers no security advantage over classical substitution ciphers and therefore no quantum-relevant one, and we state the conditions under which this negative result would be falsified.

1. Introduction

A recurring failure mode in cryptographic engineering is the deployment of primitives whose security rests on the secrecy of a design parameter rather than on the secrecy of a key under a public, well-scrutinized algorithm. The practitioners who face this problem concretely are security architects performing cryptographic migration and inventory work: the Cryptographic Bill of Materials (CBOM) methodology exists precisely because organizations must decide which deployed mechanisms to trust and which to retire as algorithms are deprecated and post-quantum readiness is planned [5]. If a mechanism such as unknown-radix encoding were believed to be a viable lightweight cipher or key-encapsulation trapdoor, it could enter that inventory as a candidate; if it is in fact breakable by classical statistical attacks on a single ciphertext, it must be excluded before it is ever deployed. Today no formal analysis exists that lets an architect, a standards body, or a protocol designer make that decision with confidence: the proposal circulates as an informal conjecture that a secret base B — the observation that a positional numeral cannot specify its own base — could act as a cryptographic trapdoor [10].

The stakes are concrete. An organization that adopted such a scheme for key encapsulation would believe its encapsulated keys were protected by a large key space — over 2^{100} combinations of radix and digit map in the parameterization analyzed here — and might further believe that a scheme built from integer arithmetic rather than from structured lattices or block ciphers resists quantum attack by being "unusual." This paper shows that both beliefs are wrong: the entire key space collapses under a statistical test that costs less than 10^8 elementary operations, and the residual problem is a monoalphabetic substitution cipher, a primitive broken by classical statistical means long before the era of digital computers. The decision that changes if this result holds is simple: unknown-radix encoding is classified as obfuscation, not encryption, and is excluded from cryptographic inventories and migration plans at the same tier as encoding tricks, with no post-quantum consideration owed to it.

This paper treats the scheme named Silent Radix Encryption in the source literature [10], which states that SRE uses a secret base b as a trapdoor for key encapsulation and that its security reduces to an unknown-weight integer knapsack problem. We analyze the core encoding primitive directly and show that the knapsack reduction, even if correct, does not confer security, because the encoding admits a direct statistical attack that bypasses the knapsack structure entirely.

Contributions.

- A formal definition of the unknown-radix encoding scheme, its key space (admissible radices $B \in [2, B_{\max}]$ and digit maps over an m -symbol alphabet), and the computation of its entropy: 104.38 bits for $m = 26$, $B_{\max} = 2^{16}$ (Sections 3 and 4.1).
- A reduction showing that decryption under a known radix is exactly a monoalphabetic substitution with base-dependent segmentation, and a derivation of the unicity distance $n_0 = 149.1$ symbols for the full key space (Sections 4.2 and 4.3).
- An explicit two-stage ciphertext-only attack with full arithmetic: radix recovery by histogram discrimination at a cost of at most 6.5535×10^7 digit operations for $n = 1000$, followed by frequency analysis of the residual substitution (Section 4.4).
- A comparison against brute-force benchmarks showing that the key-space entropy of 104.38 bits is structurally irrelevant: the binding constraint is the unicity distance, not the key count (Sections 4.5 and 5).
- A statement of the negative result for quantum relevance — a scheme broken classically on a single ciphertext cannot offer a quantum-relevant advantage — together with falsification conditions and the limits of the analysis (Section 6).

2. Background and Related Work

The immediate ancestor of this work is the silent-radix proposal itself. The QNFO note on Silent-Radix Cryptography states the core observation — that a positional numeral cannot internally specify its own base, so the base is "silent" — and proposes Silent Radix Encryption (SRE), which uses a secret base b as a trapdoor for key encapsulation and claims that security reduces to an unknown-weight integer knapsack problem [10]. A companion QNFO note connects the silent-radix principle to factoring and adelic complexity [9]; the summary supplied for that entry is empty, so we can relate it to our argument only through its title, and we do not rely on it. A third QNFO note, HM-LWE v3.0, proposes hidden-modulus LWE as a novel hardness assumption [11]; the supplied summary states only that title-level claim, so we cite it solely as evidence that novel hardness assumptions of this family are being proposed and require exactly the kind of adversarial analysis this paper performs for the radix case.

The broader literature supplies the evaluation standards against which any new primitive must be judged. Kerckhoffs's principle — the system must remain secure even if everything about it except the key is public — is operationalized in modern formal methods: the first formal model of Bitcoin's Payment Protocol (BIP70) formalised the previously shown refund-attack vulnerability arising from lack of authentication of refund addresses [4]. That work is the template we follow: an informal belief about a protocol's security is converted into a formal model, and the model is then attacked. Our Section 3 formalization and Section 4 attack play the same role for unknown-radix encoding that the formal model and refund-attack analysis played for BIP70.

On the constructive side, the literature shows what real key-establishment machinery looks like, and thereby sharpens the comparison in Section 4.5. The Multi-Factor Key Derivation Function (MFKDF) construction extends password-based key derivation with additional authentication factors such as TOTP, HOTP, and hardware tokens, and reports an exponential security improvement over PBKDFs at less than 12 ms of added latency [7]. The relevant lesson is that practical key management derives security from multiple independent secret inputs under public

algorithms — the exact opposite of deriving it from a single obscure encoding parameter. Biometric key-generation work makes the same point from another direction: two closely related papers on generating cryptographic keys from multimodal biometric feature-level fusion of fingerprint and iris address the problem that users cannot remember long keys, by deriving repeatable keys from user biometrics while incorporating the volatility of those features [2], [3]. Both accept that the key must be a high-entropy secret consumed by a standard primitive; neither proposes that the encoding of data could itself carry the secrecy. Our result explains why: encoding secrecy is structurally weak.

The threat model and deployment context are framed by work on how security mechanisms actually fail in the field. An analysis of Extended Reality (XR) devices presents a device-centric survey of security and privacy attacks and defenses and argues for robust, privacy-aware mechanisms and design considerations [8]; a study of non-expert users' perspectives on cyber security and privacy, based on online discussions on Twitter, analyzes those perspectives for traditional and smart home and IoT devices [6]; the entry's summary states no further findings, and the inference that non-expert deployers cannot evaluate mechanisms whose security depends on subtlety is this paper's own, not a finding of [6]. Both support the practical stakes in Section 7: mechanisms whose security depends on subtlety (such as a silent radix) are precisely the mechanisms non-expert deployers cannot evaluate, so a provable negative result is the only protection that scales. Finally, work on cryptographically verifiable authorization for autonomous AI agents proposes a falsifiable hypothesis and proof of concept for giving agents' requests cryptographic evidence of policy compliance in context [1]. That work illustrates the standard a new primitive must meet — a falsifiable, verifiable security claim with a proof of concept — and Section 6 states the falsification conditions for our own claim in the same spirit.

Two methodological gaps emerge from this survey. First, no cited work analyzes base-ambiguity or radix-obscure constructions; the silent-radix proposal [10] asserts a reduction to knapsack but supplies no key-space entropy, no unicity distance, and no attack cost. Second, the cited formal-analysis work [4] shows the value of converting informal proposals into models, but has not been applied to encoding-based primitives. This paper fills both gaps for the unknown-radix case.

3. Methods

3.1 The scheme

Let Σ be a plaintext alphabet of m symbols (for concreteness, English letters, $m = 26$). The secret key is a pair $K = (B, \pi)$ where:

- $B \in [2, B_{\max}]$ is the radix, an integer, with $B_{\max}$ a public parameter of the scheme;
- $\pi : \Sigma \rightarrow \{0, 1, \dots, m - 1\}$ is a secret bijection assigning each plaintext symbol a digit value.

Encryption of a plaintext $P = s_1 s_2 \dots s_n$ of n symbols proceeds in two steps:

1. Map each symbol to its digit: $d_i = \pi(s_i)$ for $i = 1, \dots, n$.
2. Interpret the digit string $d_1 d_2 \dots d_n$ as a positional numeral in base B , i.e., compute the integer

$$C = \sum_{i=1}^n d_i B^{n-i}.$$

The ciphertext is the integer C , transmitted without base information. Decryption by the holder of K writes C in base B , recovers the digit string $d_1 \dots d_n$, and applies π^{-1} .

This is the encoding core of the scheme described as Silent Radix Encryption in [10], restricted to the pure encoding primitive; we do not model any additional encapsulation layers, and our conclusions apply to the core. We note that the source proposal claims a reduction of SRE security to an unknown-weight integer knapsack problem [10]; Section 4.4 shows that the core encoding is attacked without touching any knapsack structure.

3.2 Key space and adversary model

The key space is

$$\mathcal{K} = \{(B, \pi) : B \in \{2, \dots, B_{\max}\}, \pi \in \text{Sym}(m)\},$$

where $\text{Sym}(m)$ is the set of bijections from Σ to $\{0, \dots, m-1\}$. Its size is

$$|\mathcal{K}| = (B_{\max} - 1) \cdot m!.$$

The adversary model is ciphertext-only: the attacker sees the integer C (and possibly several such integers under the same key), knows the scheme, the alphabet size m , and $B_{\max}$, but not B or π . This is the Kerckhoffs-consistent model: everything about the design except K is public. We verify the model's assumptions against the scheme's own description: the base is "silent" — carried by neither the ciphertext nor any public parameter — so its secrecy is the entire claim of the primitive [10].

3.3 Language model

Frequency analysis requires a source model. We adopt the following explicitly assumed parameters for English text over $m = 26$ symbols, used consistently throughout and labeled as assumptions wherever they enter a computation:

- A_1 : per-symbol source entropy $H_1 = 4.0$ bits (against a maximum of $\log_2 26 \approx 4.70$ bits for a uniform source), giving redundancy per symbol $D = \log_2 26 - H_1 = 0.70$ bits.
- A_2 : the most frequent symbol occurs with probability $p_{\max} = 0.12$.
- A_3 : the least frequent symbol occurring in ordinary text has probability $p_{\min} = 0.005$.

These are modeling assumptions about the plaintext source, not measurements; Section 6 states how the results shift if they are varied.

3.4 Analysis methods

We use four analytic instruments, all standard and all applied with explicit arithmetic in Section 4: (i) key-space entropy $H(K) = \log_2 |\mathcal{K}|$; (ii) the unicity distance $n_0 = H(K)/D$, the number of ciphertext symbols beyond which only one key yields a plausible plaintext; (iii) a histogram discrimination test for the radix, comparing the digit-frequency skew of the true base against the near-uniform digit frequencies of a wrong base; and (iv) brute-force cost arithmetic for benchmarking. No simulations were run and no empirical attack data were collected; every number in Section 5 is either derived in Section 4 or labeled a projection with stated assumptions.

4. Analysis

4.1 Key-space entropy

With $m = 26$ and $B_{\max} = 2^{16} = 65536$:

- Number of admissible radices: $B_{\max} - 1 = 65535$.
- Number of digit maps: $26!$.

Compute $\log_2(26!)$. We have $26! = 403291461126605635584000000 \approx 4.0329 \times 10^{26}$. Then

$$\log_2(26!) = \log_{10}(4.0329 \times 10^{26}) \times \log_2 10 = 26.6055 \times 3.32193 = 88.38 \text{ bits.}$$

Step by step: $\log_{10}(4.0329) = 0.6055$, so $\log_{10}(26!) = 26.6055$; multiplying by $\log_2 10 = 3.32193$ gives $26.6055 \times 3.32193 = 88.38$.

The radix contributes $\log_2 65535$ bits. Since $65535 = 2^{16} - 1 = 2^{16}(1 - 2^{-16})$,

$$\log_2 65535 = 16 + \log_2(1 - 2^{-16}) \approx 16 - 2^{-16} \cdot \log_2 e \approx 16 - 0.000022 = 15.99998,$$

which we carry as 16.00 bits rounded to two decimals.

Total key entropy:

$$H(K) = \log_2 [(B_{\max} - 1) \cdot 26!] = 16.00 + 88.38 = 104.38 \text{ bits.}$$

This is the number a naive analysis would stop at: the key space holds $2^{104.38} \approx 2.64 \times 10^{31}$ keys, comparable to a substantial symmetric key length. Section 4.5 shows this number is irrelevant to security.

4.2 Reduction to monoalphabetic substitution

Fix the correct radix B . Decryption is: write C in base B to obtain the digit string $d_1 \dots d_n$, then apply π^{-1} . Because each plaintext symbol s_i maps to exactly one digit $d_i = \pi(s_i)$ independent of position, the map from plaintext symbols to ciphertext digits is position-independent: the scheme with known B is exactly a monoalphabetic substitution cipher over the digit alphabet $\{0, \dots, m-1\}$, with the substitution table being π .

The only structural difference from a textbook substitution cipher is segmentation: the attacker does not know where the digit boundaries fall, because that depends on B . Once B is known, segmentation is determined. Therefore the total security of the scheme is at most the security of a monoalphabetic substitution cipher, plus whatever protection the unknown segmentation adds. The following two subsections bound both terms.

4.3 Unicity distance

The unicity distance for a cipher with key entropy $H(K)$ over a redundant source with redundancy D bits per symbol is

$$n_0 = \frac{H(K)}{D}.$$

With $H(K) = 104.38$ bits (Section 4.1) and the assumed redundancy $D = 0.70$ bits/symbol (assumption A_1):

$$n_0 = \frac{104.38}{0.70} = 149.1 \text{ symbols.}$$

Arithmetic: $104.38/0.70 = 149.114\dots$, rounded to 149.1.

For the radix alone (if π were public and only B secret):

$$n_0^{(\text{radix})} = \frac{16.00}{0.70} = 22.9 \text{ symbols.}$$

Arithmetic: $16.00/0.70 = 22.857\dots$, rounded to 22.9.

Interpretation: a single ciphertext of 150 or more symbols is, in principle, sufficient to determine the key uniquely by statistical evidence alone. The unicity distance does not give an algorithm, but it bounds from below the message length at which any ciphertext-only attack must succeed, and it shows that no parameter choice in B or π raises this bound except by increasing $H(K)$ — which, as the next subsection shows, does not help because B is recoverable directly.

4.4 The two-stage attack

Stage 1: radix recovery by histogram discrimination. The attacker guesses a candidate radix B' and computes the base- B' digit string of C , then examines the digit histogram. Two facts drive the test.

(a) Under the correct $B' = B$, the digit string is $\pi(s_1) \dots \pi(s_n)$, so digit frequencies equal plaintext letter frequencies permuted by π . In particular, the most frequent digit occurs with relative frequency $p_{\max} = 0.12$ (assumption A_2).

(b) Under a wrong $B' \neq B$, the integer C behaves as an arbitrary integer in base B' , and its digits are spread over B' values. As an explicitly labeled projection, we assume a uniformity heuristic: under a wrong $B' \neq B$ the integer C behaves as an arbitrary integer in base B' , so its n digits are approximately uniformly distributed over B' values, giving an expected most-common-digit frequency of approximately $1/B' + O(1/n)$; for $B' \geq 2^{16}$ and $n = 1000$ this is at most about $1.53 \times 10^{-5} + O(10^{-3})$, i.e., on the order of 10^{-3} at worst from discretization, versus 0.12 for the correct base. This heuristic is an assumption, not a derivation; Section 6 lists it among the limitations.

The test statistic is the maximum digit frequency $\hat{f}_{\max}$. The gap between the two hypotheses is a factor of at least

$$\frac{0.12}{10^{-3}} = 120$$

for $n = 1000$ and $B' \geq 2^{16}$, and grows linearly in n . For smaller candidate bases $B' \leq m$, the digit alphabet is compressed but the correct base still yields the permuted letter-frequency profile, which is highly non-uniform, while a wrong small base distributes the integer's value over few digits in a pattern that does not match a permutation of the letter-frequency profile; the same skew test applies.

Cost of Stage 1: for each candidate $B' \in \{2, \dots, B_{\max}\}$, extracting the n base- B' digits of C costs $O(n)$ division steps. Total:

$$W_1 \leq (B_{\max} - 1) \cdot c \cdot n$$

elementary operations, with c a small constant (we take $c = 1$ digit-extraction step per digit as the counting convention). For $B_{\max} = 2^{16}$ and $n = 1000$:

$$W_1 \leq 65535 \times 1000 = 65,535,000 = 6.5535 \times 10^7 \text{ digit operations.}$$

Arithmetic: $65535 \times 1000 = 65,535,000$.

At a conservative 10^8 elementary operations per second on a single core, this is

$$\frac{6.5535 \times 10^7}{10^8} = 0.66 \text{ s;}$$

at 10^9 /s, 0.066 s. Even granting a factor-of-100 overhead for histogram bookkeeping, Stage 1 completes in about a minute on one core: $0.66 \text{ s} \times 100 = 66 \text{ s}$ at 10^8 operations/s, and 6.6 s at 10^9 operations/s.

Stage 1b: consistency check on the digit map's range. Independently, the observed maximum digit value $d_{\max}$ in the correct segmentation satisfies $B \geq d_{\max} + 1$, and the number of distinct digits observed bounds m from below. The probability that a digit whose plaintext symbol has probability $p \geq p_{\min} = 0.005$ (assumption A_3) fails to appear in n symbols is at most

$$(1 - p_{\min})^n = 0.995^{1000} = e^{1000 \ln 0.995} = e^{-5.01254} = 0.006654.$$

Arithmetic: $\ln 0.995 = -0.00501254$; multiplied by 1000 gives -5.01254 ; and $e^{-5.01254} = 0.006654$, closing the displayed computation $(1 - p_{\min})^n = 0.995^{1000} = e^{-5.01254} = 0.006654$.

So with $n = 1000$ symbols, a symbol with frequency $p_{\min} = 0.005$ appears with probability $1 - 0.006654 = 0.993346$; the correct segmentation exhibits essentially the full letter-frequency profile, and any candidate base whose digit histogram lacks the expected skew is eliminated. The check costs $O(n)$ per surviving candidate and does not change the Stage 1 cost bound.

Stage 2: frequency analysis of the residual substitution. With B recovered, the attacker holds the digit string $d_1 \dots d_n$, which is the plaintext under a monoalphabetic substitution π . Standard frequency analysis against English letter statistics recovers π ; the unicity distance computed in Section 4.3 shows that $n \geq 149.1$ symbols suffice in principle for unique determination, and $n = 1000$ exceeds this by a factor of 6.7 ($1000/149.1 = 6.71$). No new machinery is needed: the attack is the classical one, and its cost is dominated by Stage 1 in our parameterization.

4.5 Brute-force benchmark and the irrelevance of key-space size

For contrast, consider exhaustive search over $\mathcal{K}$. With $|\mathcal{K}| = 2^{104.38} \approx 2.64 \times 10^{31}$ keys (Section 4.1) and a hypothetical 10^9 keys/s:

$$T_{\text{bf}} = \frac{2.64 \times 10^{31}}{10^9} \text{ s} = 2.64 \times 10^{22} \text{ s}.$$

Converting to years at 3.156×10^7 s/year:

$$T_{\text{bf}} = \frac{2.64 \times 10^{22}}{3.156 \times 10^7} = 8.36 \times 10^{14} \text{ years}.$$

Arithmetic: $2.64/3.156 = 0.8364$, and $10^{22}/10^7 = 10^{15}$, giving $0.8364 \times 10^{15} = 8.36 \times 10^{14}$ years.

The comparison is the paper's central quantitative point: exhaustive search needs 8.36×10^{14} years, while the structural attack of Section 4.4 needs 0.66 s of single-core work at the conservative rate of 10^8 operations/s. The ratio is

$$\frac{8.36 \times 10^{14} \text{ years}}{0.66 \text{ s} \times (1 \text{ year}/3.156 \times 10^7 \text{ s})} = \frac{8.36 \times 10^{14}}{2.09 \times 10^{-8}} \approx 4.0 \times 10^{22}.$$

Arithmetic: $0.66/3.156 \times 10^7 = 2.09 \times 10^{-8}$ years; $8.36 \times 10^{14}/2.09 \times 10^{-8} = 4.0 \times 10^{22}$. A key-space-size argument therefore overestimates the attack cost by roughly 10^{22} ; the binding constraint is the unicity distance and the statistical structure of the plaintext, not the key count.

5. Results

All numbers below are derived in Section 4 under the stated assumptions A_1 – A_3 ; none is a simulation or measurement.

- **R1 (key-space entropy).** For $m = 26$, $B_{\max} = 2^{16}$: $H(K) = 104.38$ bits, $|\mathcal{K}| \approx 2.64 \times 10^{31}$ (Section 4.1).
- **R2 (unicity distance).** $n_0 = 149.1$ symbols for the full key; $n_0^{(\text{radix})} = 22.9$ symbols if π is public (Section 4.3).
- **R3 (radix-recovery cost).** Stage 1 costs at most 6.5535×10^7 digit operations for $n = 1000$, i.e., 0.66 s at 10^8 ops/s on one core, or 0.066 s at 10^9 ops/s (Section 4.4).
- **R4 (discrimination gap, projection).** Under the uniformity heuristic stated in Section 4.4, the maximum-digit-frequency gap between correct and wrong base is a factor of at least 120 for $n = 1000$ and $B' \geq 2^{16}$ (Section 4.4); this is a projection conditional on that heuristic, not a derived bound.
- **R5 (appearance probability).** A symbol of frequency $p_{\min} = 0.005$ fails to appear in $n = 1000$ symbols with probability 0.006654 (Section 4.4, Stage 1b).
- **R6 (brute-force benchmark).** Exhaustive search at 10^9 keys/s takes 8.36×10^{14} years; the structural attack is faster by a factor of about 4.0×10^{22} (Section 4.5).
- **R7 (negative result).** With B recoverable at cost R3 and the residual cipher a monoalphabetic substitution breakable within R2 symbols, the scheme's ciphertext-only security on a single message is bounded by the cost of classical frequency analysis; radix obscurity contributes no security beyond obfuscation, and hence no quantum-relevant advantage (Sections 4.2–4.5).

6. Discussion

The result is a negative one, and its strength depends entirely on the assumptions behind the arithmetic. We state the limitations, the failure modes, and the falsification conditions explicitly.

Limitations. The result rests on the following assumptions, with the sensitivity of the headline numbers to each:

- *Language model (A_1 – A_3).* The redundancy $D = 0.70$ bits/symbol sets the unicity distance $n_0 = H(K)/D$; if the true redundancy were $D = 0.5$, then $n_0 = 104.38/0.5 = 208.8$ symbols, and if $D = 1.0$ (near-uniform source), $n_0 = 104.4$ symbols. The qualitative conclusion — a single message of a few hundred symbols suffices in principle — is stable across this range, but the exact threshold scales as $1/D$.
- *Frequency skew (A_2).* The Stage 1 gap factor of 120 scales linearly with $p_{\max}$; a flatter source with $p_{\max} = 0.05$ would reduce the gap to $0.05/10^{-3} = 50$, still decisive for $n = 1000$.
- *Cost convention.* The Stage 1 cost takes $c = 1$ operation per digit extraction; real implementations pay a larger constant, which is why we granted a factor-of-100 overhead and still obtained sub-minute runtimes. The bound is on digit operations, not wall-clock time on any specific machine.
- *Scope of the model.* We analyze the pure encoding core of the scheme described in [10]. If the full SRE construction adds layers that randomize digit frequencies or re-encrypt per message, the Stage 1 histogram test may fail; the source proposal's supplied description states only the trapdoor claim and the knapsack reduction, and gives no such layer, so we cannot assess one.
- *Ciphertext-only model.* The attack assumes the attacker sees only C . A known-plaintext adversary recovers π directly from a single aligned pair and B from the digit range, making the scheme trivially broken; our model is the weaker and therefore stronger setting.

Failure modes of the analysis. The reduction in Section 4.2 assumes the digit map is position-independent, which follows from the encoding definition; if a variant made π position-dependent (a polyalphabetic generalization), the residual cipher after Stage 1 would be a polyalphabetic substitution, and the unicity distance would change because

the effective key entropy per message would grow. Our result does not cover such variants.

Falsification. The main result — that radix obscurity adds no security beyond obfuscation — would be falsified by a concrete demonstration of either of the following: (i) a parameterization of the scheme in which the digit histogram under the correct base is statistically indistinguishable from that under all wrong bases for message lengths at which the scheme is claimed to be usable, i.e., a failure of the Stage 1 discrimination gap of Section 4.4; or (ii) a variant construction whose security proof does not pass through the encoding core analyzed here, for example one where the transmitted object is not the integer C but a transformed value that destroys the base- B digit structure. Absent either, the negative result stands: any scheme whose ciphertext is the bare positional integer in a secret base falls to the two-stage attack.

Open questions. Whether the knapsack reduction claimed in [10] is correct is not settled here, and it does not need to be: even a correct reduction to an NP-hard problem would not help, because the direct attack bypasses the knapsack structure. Whether some non-encoding layer built on top of the core could be secure is a separate design question, not a rescue of the radix-as-trapdoor idea.

7. Practical Implications and Engineering Pathways

(a) Who uses the result and what changes. The direct users are security architects performing cryptographic inventory and migration work under the CBOM methodology [5], and standards reviewers evaluating novel primitive proposals. If this analysis holds, the decision changes as follows: unknown-radix encoding is classified as obfuscation rather than encryption, is recorded in a cryptographic bill of materials at the encoding tier rather than the cipher tier, and is excluded from post-quantum migration planning, because a scheme broken classically on a single ciphertext requires no quantum consideration. The proposal's own framing — a trapdoor for key encapsulation [10] — would be rejected at the design-review stage on the strength of the Stage 1 cost bound alone.

(b) Readiness level. This work is purely analytical: a formalization, entropy computation, and attack derivation with no implementation, no simulation, and no peer review. Against the NASA Technology Readiness Level definitions this corresponds to TRL 1–2 (basic principles observed and formulated; concept and application formulated). Analytical work does not exceed TRL 3, and no claim of implementation or manufacturing readiness is made or possible here.

(c) Smallest next experiment. The minimal de-risking (or falsifying) test is a software prototype implementing the encoding of Section 3.1 for $m = 26$ and a spread of radices $B \in \{2^4, 2^8, 2^{16}\}$, encrypting English texts of $n \in \{150, 500, 1000\}$ symbols, and running the Stage 1 histogram test to measure the empirical radix-recovery success rate. The outcome that refutes the paper's claim is a radix-recovery rate significantly below 1 at $n = 1000$ — specifically, if the correct base is not ranked first by the maximum-digit-frequency statistic in a majority of trials, the discrimination gap of Section 4.4 fails and the negative result must be revised.

(d) Binding constraints. For the prototype test itself, the binding constraints are modest: the arithmetic of Section 4.4 shows the attack needs on the order of 10^8 digit operations, well within a single-core software implementation, so compute is not a constraint; the constraint that binds first is the fidelity of the plaintext corpus to the language model A_1 – A_3 , since a corpus with unrepresentative letter frequencies would distort the measured discrimination gap. For any hypothetical deployment of the scheme, the binding constraint is the one this paper establishes: the scheme cannot be deployed as encryption at all, so materials, tolerances, and scale-up constraints never arise. Regulatory constraints enter only through the CBOM classification decision in (a).

(e) Order-of-magnitude bounds. All performance figures in this paper are derived in Section 4 from stated inputs: the Stage 1 bound of 6.5535×10^7 digit operations assumes $c = 1$ operation per digit extraction and $n = 1000$; the wall-clock figures of 0.66 s and 0.066 s assume 10^8 and 10^9 elementary operations per second

respectively on a single core. These are projections from the stated operation counts and rates, not measurements; a factor-of-100 bookkeeping overhead was granted and does not change the sub-minute conclusion.

(f) Risks and misuse paths. The misuse risk runs in the opposite direction from most cryptographic papers: the danger is not that attackers exploit the scheme, but that defenders deploy it. A non-expert implementer, encountering the "silent radix" observation, might adopt the encoding as a lightweight cipher in a low-resource device on the mistaken belief that a 2^{104} key space confers security [10]. The study of non-expert users' perspectives on cyber security and privacy for traditional and smart home and IoT devices [6] shows that non-expert perspectives on such devices are a documented subject of analysis; that non-expert implementers may misjudge such mechanisms is this paper's inference, and it makes the deployment path a realistic concern; the mitigation is exactly the classification decision in (a), published as an explicit negative result.

8. Conclusion

We formalized the unknown-radix encoding at the core of the Silent Radix Encryption proposal, computed its key-space entropy as 104.38 bits for a 26-symbol alphabet and radices up to 2^{16} , and showed that this entropy is security-irrelevant. The scheme with a known radix is exactly a monoalphabetic substitution cipher; the radix itself is recoverable by a histogram test costing at most 6.5535×10^7 digit operations for a 1000-symbol ciphertext — about 0.66 s on a single core at a conservative operation rate — while exhaustive search over the same key space would take 8.36×10^{14} years. The unicity distance of 149.1 symbols, computed from the key entropy and an assumed source redundancy of 0.70 bits/symbol, shows that a single short message suffices for unique key determination in principle. The conclusion for practice is categorical: a positional numeral's silent base is not a secret in any cryptographic sense, unknown-radix encoding is obfuscation and not encryption, and it offers no quantum-relevant advantage because it offers no classical security to begin with. The result would be falsified by a parameterization whose digit histograms are statistically indistinguishable across bases at usable message lengths, or by a variant construction that does not transmit the bare positional integer; we have stated both conditions so that the claim is testable.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.16.0-revise-retry). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; statements about a cited work are restricted to its retrieved title and abstract.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond its abstract; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Cryptographically verifiable authorization for autonomous AI agents: a falsifiable hypothesis and proof of concept. arXiv:2607.21325v3. <https://arxiv.org/abs/2607.21325v3>
- [2] Secured Cryptographic Key Generation From Multimodal Biometrics: Feature Level Fusion of Fingerprint and Iris. arXiv:1003.1458v1. <https://arxiv.org/abs/1003.1458v1>
- [3] Secured Cryptographic Key Generation From Multimodal Biometrics Feature Level Fusion Of Fingerprint And Iris. arXiv:1002.2527v1. <https://arxiv.org/abs/1002.2527v1>
- [4] Formal Modelling and Security Analysis of Bitcoin's Payment Protocol. arXiv:2103.08436v1. <https://arxiv.org/abs/2103.08436v1>
- [5] Architecture-Derived CBOMs for Cryptographic Migration: A Security-Aware Architecture Tradeoff Method. arXiv:2603.22442v2. <https://arxiv.org/abs/2603.22442v2>
- [6] Perspectives of Non-Expert Users on Cyber Security and Privacy: An Analysis of Online Discussions on Twitter. arXiv:2206.02156v2. <https://arxiv.org/abs/2206.02156v2>
- [7] Multi-Factor Key Derivation Function (MFKDF) for Fast, Flexible, Secure, & Practical Key Management. arXiv:2208.05586v3. <https://arxiv.org/abs/2208.05586v3>
- [8] Augmenting Security and Privacy in the Virtual Realm: An Analysis of Extended Reality Devices. arXiv:2402.03114v1. <https://arxiv.org/abs/2402.03114v1>
- [9] QNFO: FACTORING, Adelic Complexity, and the Silent-Radix Principle
- [10] QNFO: Silent-Radix Cryptography: Exploiting the Base Ambiguity of Positional Notation as a Cryptographic Primitive
- [11] QNFO: HM-LWE v3.0