ENI6MA Cypher: Toward Perfect Secrecy by Design

by Frank Dylan Rosario and Dr. Lin Wang

One-Time Projective Ring-Rotation Diffusion and Private-Map Confusion in Interactive Proofs of Knowledge

Abstract

We presnt a formal, end-to-end security argument for an ENI6MA-style interactive proof in which the attacker "gains nothing," in the precise sense of Shannon's perfect secrecy. We analyze a mechanism whose central features are: **(i)** a one-time, entropy-dependent projection ("OTP") of the alphabet into six perceptual leaves; **(ii) ring-rotation diffusion** that spreads symbol appearances across a large configuration space with multiplicity $x = 30 \times 30 \times 12 = 10{,}800$ per ring-state; **(iii) confusion** induced by a **private bijective map** from the six leaves to six private synonyms (e.g., six bearing directions); and **(iv)** a **second-layer map** enumerating all permutations of the six leaves, i.e., $6! = 720$. We also quantify the combinatorial blow-up attributed to these layers by showing how the second-layer permutations amplify the ring-state space to $x^{6!} = x^{720}$ candidate configurations in a natural upper-bound model of the attacker's uncertainty. When folded into an interactive ceremony whose **responses are masked** (e.g., via XOR or modular masking) and whose **probes are balanced**, the observable transcript becomes strictly **non-informative** about the pre-commitment alphabet: **every observable value is equally probable**, across rounds and across epochs.

The first half of the work is information-theoretic. We formalize perfect secrecy for the observable transcript C, the hidden message M (the next alphabetic symbol in the commitment), the one-time projection P, the ring-rotation state R, and the private map K. Under mild regularity (uniformity and independence of the one-time elements, correct nonce/epoch discipline, and balanced probes), we prove $I(M;C) = 0$ and give explicit statements such as $P(M = m \mid C = c) = P(M = m)$ for all m, c. Intuitively, the transcript is a one-time pad in the *alphabetic geometry* (rather than in the bit strings) because a uniformly random, single-use projection and mask render the observed synonym statistically independent of the true symbol and of the private map.

The second half is computational and operational. We show that even a full-power adversary, classical or quantum, does not find traction because **no residue** (no bias, no correlations, no frequency signature) accumulates across sessions. The best the adversary can attempt is brute-forcing latent states jointly: the private map (size 6!) and the ring-rotation configuration (size x) raised to however many independent latent degrees of freedom govern the mapping. In the natural upper-bound model where each factorial degree can independently carry a ring-state, the combined space is $x^{6!} = 10800^{720} \approx 10^{2904} \approx 2^{9647}$ possibilities. Grover acceleration (square-root speedup) against such a space leaves $\approx 2^{4823.5}$; that number remains beyond reach by any remotely plausible quantum hardware. In the practical model where the ring-state renews per round (say h rounds), the attacker still faces $6! \cdot x^h$ possibilities while the transcript remains equiprobable and non-informative; no post-processing of an uninformative transcript produces information.

The human/agent solving efficiency is explained by **Gestalt search** over six perceptually salient leaves and by **XOR/modular masking** that cancels the verifier's challenge while never disclosing the raw witness. Humans locate the target alphabet quickly by preattentive features; agents compute the same mapping deterministically. The speed enjoyed by legitimate provers stems from *asymmetry of knowledge* (they know K and the current symbol) rather than from any structure available to the attacker: this is the essence of "knowing without showing."

---

1. Introduction and Motivation

Shannon's perfect secrecy formalizes a gold standard for confidentiality: observing ciphertext tells you **nothing** about the message. In symbols, **perfect secrecy** is the condition

$$P(M = m \mid C = c) \;=\; P(M = m) \quad \text{for all } m, c,$$

"The probability of message m given ciphertext c equals the prior probability of m."

Equivalently, the mutual information between M and C is zero, $I(M; C) = 0$. In conventional cryptography, the one-time pad achieves this when the key is uniform, at least as long as the message, and used once.

This dissertation considers a different, but mathematically analogous, setting: an **interactive proof** in which a prover demonstrates knowledge of the next symbol of a pre-committed secret **without revealing that symbol**. Instead of XOR'ing a bitstring with a one-time key, we **project the alphabet** into six perceptual leaves using a **one-time, entropy-dependent projection** P, we **diffuse** symbol locations with **ring rotations** across a large state space $x = 10{,}800$, and we hide the relationship between leaves and the prover's private synonyms via a **private bijection** K over the six leaves ($|K| = 6! = 720$).

Finally, we combine these with a **masked response rule** (e.g., XOR or modular addition on the index) so that the verifier can check correctness **without** learning the raw witness. The end effect is a **Shannon-style OTP in the alphabetic geometry**: the observable symbol is independent of the true symbol because of the *combined* one-time projection and mask; frequency, correlation, and cross-session inference receive no foothold.

The two reasons an attacker learns nothing are:

1. **Information-theoretic independence.** The observable transcripts are generated by a one-time random projection and masking that render them statistically independent of the message. There is no amount of computation, classical or quantum, that can turn independent data into dependent data.

2. **Combinatorial explosion in hidden state.** If an attacker insists on brute force guessing of hidden state rather than relying on information from transcripts, the attack crater is enormous. The private map space contributes a 6! factor; the ring-rotation diffusion contributes x states per applicable degree of freedom; the natural upper bound is $x^{6!}$ when rotations bind independently to each factorial component. This is so large that even Grover's quadratic speedup does not change the feasibility conclusion.

We proceed by formalizing the objects, proving perfect secrecy, quantifying spaces, and contrasting the **efficient** human/agent solve with the **intractable** attacker search. Throughout, we keep to two reading tracks: (i) rigor with equations; and (ii) short **TTS** paraphrases after each displayed formula for accessible narration.

---

2. Objects and Ceremony: Symbols, Projections, Maps, and Masks

Let the **message space** M denote the set of possible next alphabetic symbols (e.g., 62 symbols for upper, lower, digits, or any application-specific alphabet). Let $M \in M$ denote the prover's next true symbol.

Let there be **six perceptual leaves** (color/shape/orientation zones), indexed by $L = 0, 1, 2, 3, 4, 5$. A **one-time projection** P is a randomized mapping that, for the epoch (round window) in question, disposes elements of M into the six leaves so that each leaf receives an **equiprobable** fraction of the alphabet, and importantly, the projection P is **one-time** and **entropy-dependent** (seeded from fresh randomness). Formally, we can model P as a draw from a family P of balanced allocations:

$$P \sim \text{Uniform}(P), \quad P : M \to L, \quad \forall \ell \in L : \ |\{m \in M : P(m) = \ell\}| \approx |M|/6.$$

"We choose a random projection that maps the alphabet to six leaves with equal load."

Let R denote the **ring-rotation diffusion state** selected uniformly from a configuration space with cardinality

$$x \;=\; 30 \times 30 \times 12 \;=\; 10{,}800.$$

"The ring state has ten thousand eight hundred possibilities."

R governs how symbol instances (glyphs, tiles, or positions) drift across the canvas via ring-like shifts, think concentric rings with 30 angular slots, 30 radial slots, and 12 ring layers. R is freshly sampled per epoch (and, if desired, per round), and its effect is that even if the attacker knew P, they do not know **where** instances occur without R.

Let K denote the **private map**, a hidden bijection $K : L \to S$, where S is a set of six **private synonyms** (e.g., $[up, down, left, right, forward, back]$). The size of the K-space is

$$|K| = 6! = 720.$$

"There are seven hundred twenty private maps."

During a round, the verifier issues a **challenge mask** $Z \in L$ (or a one-hot vector mask over the six leaves), and the prover emits a **masked response** $Y \in L$ that the verifier can check without seeing the raw leaf index $S = P(M)$. Two natural masking rules are:

- **Modular addition on indices:** $Y = S \oplus Z$ with $\oplus$ meaning addition modulo 6.
- **XOR on one-hot encodings:** encode S and Z as one-hot vectors in $0, 1^6$ and set $Y = S \oplus Z$ (bitwise XOR).

We will write the index-space version,

$$Y \;=\; S \oplus Z, \quad S = P(M) \in L.$$

"The observed response equals the secret leaf index plus the mask modulo six."

Given Z, the verifier can **cancel** Z and check whether $(Y \ominus Z)$ matches the asserted leaf index implied by the prover's private map and the verifier's ephemeral placement rules. What the verifier **cannot** learn is K or M; they only learn **consistency** across rounds.

The **transcript** C consists of the public capsule metadata for the epoch (including the hash of P's seed and any R metadata, not the raw values themselves), the sequence of masked responses $Y_1, \ldots, Y_h$, the challenges $Z_1, \ldots, Z_h$, and the final accept bit.

---

3. Shannon Perfect Secrecy in the Alphabetic Geometry

3.1 Statement of perfect secrecy

We aim to show that under correct operation (fresh P, fresh R, balanced probes/masks), the transcript C leaks nothing about M. A minimal statement is:

$$P(M = m \mid C = c) \;=\; P(M = m) \quad \text{for all } m \in M, c \in C,$$

"The posterior of the message given the transcript equals the prior."

Because C is determined by $(Y_1, \ldots, Y_h)$ and public challenge values, it suffices to show that, per round, Y is **independent** of M. If for each round i, $Y_i \perp M$, and the public metadata merely commits to the *existence* of one-time seeds without exposing them, then the whole transcript is independent.

3.2 Why one-time projection + mask force independence

Fix a round and suppress the subscript i. Let $S = P(M)$ denote the secret leaf index. If P is uniform and **fresh for the epoch**, then for any fixed m we have:

$$P\big(S = \ell \mid M = m\big) \;=\; \tfrac{1}{6}, \quad \forall \ell \in L.$$

"Conditioned on any message, the secret leaf is uniform over six choices."

If Z is an **independent** uniform mask in L and $Y = S \oplus Z$, then for each m and each observed y:

$$P(Y = y \mid M = m) \;=\; \sum_{\ell \in L} P(S = \ell \mid M = m) P(Z = y \ominus \ell) \;=\; \sum_{\ell} \tfrac{1}{6} \cdot \tfrac{1}{6} \;=\; \tfrac{1}{6}.$$

"With a uniform secret leaf and a uniform mask, the observed response is also uniform."

Thus Y is uniform and **independent** of M. This is the exact algebra that makes the classical one-time pad perfectly secret: a uniform message masked with a uniform independent key yields a uniform ciphertext independent of the message. Here the "message" that reaches the observable channel is the **leaf index** $S = P(M)$; P plays the role of a per-epoch randomization that equalizes the distribution of S given M, and Z is the per-round mask that removes any remaining structure.

Formally,

$$I(M; Y) \;=\; 0 \quad \text{and} \quad I(M; C) = 0.$$

"The mutual information between the message and the observed response, and between the message and the whole transcript, is zero."

3.3 Role of diffusion R

The ring-rotation diffusion state R does not need to carry enormous entropy per se to ensure perfect secrecy, P and Z already suffice for $I(M;Y) = 0$, but R supplies essential **operational robustness**:

1. **No spatial/frequency residue.** Diffusion spreads the whereabouts of symbol instances across x possible canvas dispositions so that even a powerful observer who tries to correlate *positions* across rounds sees no stable layout.

2. **No cross-epoch carryover.** Because R is fresh together with P, positions and local neighborhoods in one epoch are stochastically unrelated to those in the next.

Quantitatively, R sampled from $x = 10{,}800$ provides $\log_2(x) \approx 13.40$ bits of additional randomness per independent use. While 13.40 bits is modest on its own, its value compounds wherever R influences multiple independent latent placements, see §6 for how this multiplies.

---

4. Confusion by a Private Map and the Second-Layer Search Space

4.1 Private bijection over six leaves

A central ingredient is the **private map** $K \in S_6$, a permutation of six leaves into six synonyms. The "confusion" property is succinct: the **observable** symbol is a synonym label, not the leaf's public identity. Without knowledge of K, the same observed synonym can mean different public leaves across provers; and within a single prover, it associates to different public leaves across epochs due to new P and R. The cardinality is

$$|S_6| \;=\; 6! \;=\; 720, \quad H(K) \;=\; \log_2(6!) \approx 9.492 \text{ bits}.$$

"There are seven hundred twenty possible private maps, giving about nine and a half bits of entropy."

Although $H(K)$ alone is small, it composes with other randomness sources. The attacker cannot infer K from transcripts because the transcripts are independent of M and therefore do not carry any consistent alignment signal; any attempt to fit K becomes a search through an *underdetermined* model fed with independent noise.

4.2 Second-layer combinatorics with ring-rotation diffusion

The ring-rotation diffusion space $x = 10{,}800$ multiplies when attached to independent latent degrees. Two models are relevant:

1. **Per-round model (conservative and realistic).** If the ring-rotation state refreshes per round, say h rounds per proof, then the joint hidden state per proof has size at least $|K| \cdot x^h = 720 \cdot 10800^h$.

2. **Upper-bound factorial model (pessimistic to the attacker).** If the diffusion state can be treated as independently associated with each private-map degree, i.e., x choices *per permutation component*, then the hidden state explodes to $x^{6!}$, i.e.,

$$x^{6!} \;=\; 10800^{720}.$$

"The upper-bound space equals ten thousand eight hundred raised to the seven hundred twentieth power."

The latter is intended as a **design-space upper bound** the attacker must contemplate when they cannot tell *which* diffusion choices bind to *which* private-map degrees. Its size is extreme:

$$\log_{10}\!\left(10800^{720}\right) = 720\log_{10}(10800) \approx 720 \times 4.033423755 \approx 2904.065,$$

so

$$10800^{720} \;\approx\; 10^{2904.065} \;\approx\; 2^{9647.095}.$$

"The space is about ten to the two-thousand nine hundred four, equivalently around two to the nine-thousand six hundred forty-seven."

We emphasize: the **core secrecy claim** (no information leakage) does **not** rely on large search spaces; it is **information-theoretic** given fresh P and Z. The $x^{6!}$ factor matters for **brute-force fantasies** when an attacker dreams of enumerating hidden states rather than learning from transcripts.

---

5. All Values Equally Probable: No Frequency, No Correlation, No Residue

5.1 Equiprobability of observed responses

From §3, for each round and any $m \in M$,

$$P(Y = y \mid M = m) \;=\; \tfrac{1}{6} \quad \text{for all } y \in L.$$

"Whatever the message is, each observed response is equally likely, one sixth."

Therefore, any empirical frequency table the attacker builds for Y is **flat**; the law of large numbers converges the table to uniform. There is no class-conditional bias to learn, no signature to correlate across rounds or across sessions, and no residual statistics tied to M or K. This defeats **frequency analysis**, **n-gram analysis**, and any **co-occurrence** mining: the observable layer is engineered to be *statistically featureless* about the hidden values.

5.2 Equiprobability of positions and placements

Because R renews and re-disposes symbol instances over the ringed canvas, any attempt to build spatiotemporal correlations across sessions collapses. The attacker cannot exploit, say, "symbol A tends to appear near radius 12 at angle 17" because the joint distribution of positions changes with R. With balanced P and independent R, the spatial marginals are uniform within their design tolerances; across epochs, they are independent.

5.3 Why quantum does not change equiprobability

A quantum computer can speed up **search**, but it cannot **manufacture information** where none exists. If $I(M;C) = 0$, then for any algorithm (classical or quantum), the posterior remains the prior. No amplitude-amplification can amplify a **nonexistent bias**. This is the deepest reason "quantum gains nothing" in this design: the secrecy is **information-theoretic at the interface**, not computational.

---

6. Complexity Landscape for Attackers

6.1 Baseline: Information-theoretic futility

If an attacker's workflow is "observe C; infer M," perfect secrecy says there is no better inference than the prior. More formally, for any estimator $\hat{M}(C)$,

$$\Pr[\hat{M}(C) = M] \;=\; \max_{m} P(M = m) \quad \text{(no a-posteriori improvement)}.$$

"The best guess after seeing the transcript is still just the prior best guess."

The attacker may instead target **hidden state**: K, R, and any latent seeds for P. But C remains non-informative about these as well under the equiprobability conditions; hence the attacker receives **no a-posteriori narrowing** of the search space from data.

6.2 Brute force over private maps and ring states

Let us tabulate two search models.

Model A (per-round). There are $|K| = 720$ private maps and $x = 10800$ ring states per round, assumed independent across h rounds. The space is

$$S_A \;=\; 720 \cdot 10800^h, \quad \log_2 S_A \;=\; \log_2 720 + h \log_2 10800 \;\approx\; 9.492 + 13.399h.$$

"In the per-round model, the search space has seven hundred twenty times ten thousand eight hundred to the power h possibilities."

Model B (upper-bound factorial). Each of the 6! private-map degrees may be associated with an independent ring-state, this reflects a coupling pattern hopelessly opaque to the attacker. The space is

$$S_B \;=\; 10800^{720}, \quad \log_2 S_B \;\approx\; 9647.095 \text{ bits.}$$

"In the factorial upper bound, the space is about nine thousand six hundred forty-seven bits."

Even **Model A** grows quickly: with $h = 32$ rounds (a short interactive proof), $\log_2 S_A \approx 9.492 + 32 \times 13.399 \approx 438.3$ bits, already too large to search, particularly because the transcript does not provide a correctness oracle beyond the single accept bit at the end. **Model B** is astronomically larger.

6.3 Quantum search: Grover bounds do not help enough

Grover's algorithm reduces the expected number of queries from N to $O(\sqrt{N})$ for unstructured search. Applying this to the spaces above:

- For S_A with, say, $h = 32$, a $\sim 2^{438}$ space becomes $\sim 2^{219}$, still out of reach.
- For $S_B \approx 2^{9647}$, Grover leaves $\sim 2^{4823.5}$, even more fantastical.

Crucially, **Grover presumes an oracle** that identifies correct candidates. In our setting, the only oracle is a *live interaction* with the legitimate verifier or a perfect simulator of the one-time projection and mask sequence, both unavailable to the attacker without the very secrets they are trying to discover. Without an oracle, Grover's algorithm does not instantiate.

6.4 No algebraic structure for Shor or number-theoretic attacks

Shor's algorithm targets discrete logarithms and factoring; hidden subgroup methods require group structure. Our construction intentionally **exposes no algebraic trapdoor** at the transcript layer: it is PRF-level masking on a small alphabet with one-time randomization, not a modular exponentiation with a known modulus. There is nothing to diagonalize or period-find.

---

7. Human and Agent Efficiency: Why Solving Is Fast for the Right Party

7.1 Gestalt search over six leaves

Humans are adept at **preattentive feature detection**: color, orientation, and shape pop-out effects enable near-instant localization of a target cluster

among distractors. The six-leaf partition capitalizes on this: each leaf is visually distinct; the **target alphabetic class** (e.g., uppercase, lowercase, digits, or any custom grouping) is clearly associated with one leaf by the **private map** K. The human solver's loop per round is:

1. **Orient to the six leaves.** Each leaf has a distinctive gestalt; the human registers them subconsciously.
2. **Recall K.** The human remembers "for me, 'up' means the leaf that currently carries lowercase (say)."
3. **Locate the symbol cluster.** The ring-rotated positions do not hinder recognition because the *category* lives on the leaf; the exact pixel coordinates are immaterial.
4. **Mask response.** Given the challenge Z, the human computes $Y = S \oplus Z$ (six-way modular shift) by mental arithmetic or a mnemonic, and responds with the synonym corresponding to Y.
5. **Repeat.** Consistency across rounds yields acceptance.

Every sub-step is constant-time with practiced use; the difficulty is bounded and does **not** scale with the magnitude of the attacker's search space. This is the moral asymmetry: **the right party knows where to look; the wrong party sees white noise.**

7.2 Agent (AI) execution

An agent (software prover) performs the same steps deterministically:

- Compute $S = P(M)$ from the one-time projection seed.
- Compute $Y = S \oplus Z$.
- Emit the synonym label associated with index Y under K.

The runtime is $O(1)$ per round given local access to the seeds. Memory is $O(1)$ because K is only six elements and the projection seed is a fixed-size value. Verification is similarly constant-time per round. Thus, **efficiency** is not at odds with **secrecy**: perfect secrecy is achieved by design choices at the interface, not by computational heaviness.

---

8. XOR (or Modular) Masking: The Algebra That Eliminates Leakage

The algebra of masking is the fulcrum of the "no residue" property. Let $S \in 0, \ldots, 5$ be the secret leaf index, Z the mask, and $Y = S \oplus Z$. The attacker sees Y and Z but not S. If Z is uniform and independent, then for any s and y,

$$P(Y = y \mid S = s) = P(Z = y \ominus s) = \tfrac{1}{6},$$

"Conditioned on the secret index, each observed response is still one sixth likely."

Hence Y is independent of S and therefore of M. The rule generalizes immediately to XOR with one-hot encodings; we keep modular indices for clarity.

A frequent misunderstanding is: "But the attacker sees many Y's; surely something leaks?" No. If the Z's are fresh and independent per round and the projection P is fresh per epoch, the **joint** distribution of $(Y_1, \ldots, Y_h)$ remains the product of six-way uniforms. There is **no statistic** the attacker can compute whose expectation differs across messages, and therefore **no hypothesis test** with power above random guessing.

---

9. Perfect Secrecy Proof Sketch with Formal Conditioning

For completeness, we formalize the independence claim with conditioning on the latent variables. Let C denote the transcript of a single round; extend to h rounds by independence.

Let the latent tuple be (P, R, K, Z) and the observed be Y. We wish to show $I(M; Y) = 0$. Because Z is public yet fresh and because P and R are one-time latent values whose seeds are **not** revealed, we condition on Z only (the rest are integrated out).

$$\begin{aligned}
P(Y = y \mid M = m) &= \sum_{p \in P} \sum_{r \in R} \sum_{k \in S_6} P(Y = y \mid M = m, P = p, R = r, K = k)\, P(p, r, k) \\
&= \sum_{p,r,k} P\big(S \oplus Z = y \mid S = p(m)\big)\, P(p) P(r) P(k) \\
&= \sum_{p,r,k} P\big(Z = y \ominus p(m)\big)\, P(p) P(r) P(k) \\
&= \sum_{p} \Big[\tfrac{1}{6}\Big]\, P(p) \quad \text{(since } Z \text{ uniform, } r, k \text{ irrelevant)} \\
&= \tfrac{1}{6}.
\end{aligned}$$

"After summing over the one-time projection, ring state, and private map, the probability of any observed response is one sixth."

A symmetric derivation shows $P(M = m \mid Y = y) = P(M = m)$, and therefore $I(M; Y) = 0$.

---

10. Practical Security: "No Residue" in the Transcript and Beyond

The above is a *mathematical* statement about the observable Y's. In practice, adversaries try more: timing analysis, interaction patterns, or cross-context linkages. The design addresses them:

- **Timing and UI jitter.** Quantizing response windows and adding UI jitter prevents timing distributions from becoming user-specific fingerprints.
- **Session isolation.** Each epoch carries fresh (P, R), and identifiers bind to the epoch only; there is no stable cross-epoch handle.
- **Balanced probes.** The verifier ensures that the distribution of challenges Z is uniform and independent across rounds; this prevents adversarial choice of Z sequences that bias the observed Y's.

With these, the transcript remains featureless with respect to M and K. The attacker's only avenues are off-channel (e.g., exploiting implementation bugs), which are orthogonal to the cryptographic core and addressable by standard hardening.

---

11. Interpreting x^{720} and Other Spaces: Clarity on Bounds

A careful reader may ask: "Is x^{720} *really* the search space?" The right answer is **twofold**:

1. **For information leakage:** The size of any hidden space is irrelevant, perfect secrecy has already said transcripts leak nothing. Even if $x = 2$ and $6! = 1$, the distributional independence guarantees $I(M; C) = 0$ as long as the one-time properties hold.
2. **For brute-force enumeration:** When an attacker cannot learn from data, they sometimes imagine enumerating latent states; then **bounds matter**. The per-round model yields $720 \cdot 10800^h$. The upper-bound factorial model yields 10800^{720}. The latter is conceptually appropriate when **multiple independent ring-rotation choices** attach to the factorial degrees of the private map, and the attacker cannot disambiguate *how* these attach. It is explicitly an *upper bound*: the mere *existence* of such an attachment mechanism forces the attacker's uncertainty to grow multiplicatively with the number of independent attachments.

Either way, **no plausible computation**, classical or quantum, can traverse such spaces in the millisecond interaction windows that define the ceremony. And none of this helps create information where the transcript has none.

---

12. Worked Numbers and Sanity Checks

12.1 Bits of entropy

- $H(K) = \log_2(6!) \approx 9.492$ bits.
- $H(R) = \log_2(10800) \approx 13.399$ bits per independent diffusion choice.
- If R refreshes per round with $h = 24$ rounds, $H(R^{(h)}) \approx 24 \times 13.399 \approx 321.6$ bits in the per-round model.

12.2 Upper-bound crater

- $H\left(x^{6!}\right) = 720 \log_2(x) \approx 720 \times 13.399 \approx 9647.1$ bits.
- Groverized cost: $\approx 2^{4823.5}$, still intractable.

12.3 False acceptance (soundness) under random guessing

Let h be the number of rounds and suppose strict acceptance. A random impersonator (no K, no M) has per-round success probability 1/6. The chance of fluking all rounds is

$$P_{\text{FA}} = \left(\frac{1}{6}\right)^h.$$

"False accept equals one sixth to the power h."

With $h = 10$, $P_{\text{FA}} \approx 1.65 \times 10^{-8}$; with $h = 20$, $P_{\text{FA}} \approx 2.7 \times 10^{-16}$. Threshold-acceptance variants (allowing a slip or two) are governed by binomial tails, which remain negligible for modest h.

These figures are orthogonal to secrecy but underscore that **usability** (few rounds) and **security** (tiny impersonation risk) coexist comfortably.

---

13. Why the Human Mind Is Fast (and the Attacker Is Not)

Humans recognize categories in clutter **without serial enumeration**. The six-leaf layout amplifies preattentive cues: a glance suffices to decide "my category is in that leaf." The **XOR/mod-6 masking** is trivial mental arithmetic; the human never needs to remember a large table, only a six-way permutation K and a "rotate by Z" instruction. This is **fast**, **repeatable**, and **error-tolerant**.

Contrast the attacker: with no access to K and no stable mapping from the observed synonyms to public leaves (because P is one-time and R diffuses placements), the attacker's best attempt is exhaustive search over K and R (and possibly P seeds). There is **no gradient** to climb and **no statistic** to exploit.

This is what "no residue" means: after the ceremony, the attacker's posterior is the prior.

---

14. Quantum Adversaries: No Advantage Without Structure or Signal

Quantum algorithms excel when either:

- There is **structure** (periodicity, group homomorphisms) to exploit (e.g., Shor).
- There is an **oracle** that identifies winning states so amplitude amplification finds a needle faster (Grover).

Our design offers neither. The observable layer is **structureless** (equiprobable outputs), and there is no **oracle** available to the attacker to verify guesses about K or R off-line. In an on-line interaction, the verifier does not aid an attacker; and even if an attacker queried a verifier, the queries are rate-limited, audited, and tied to **one-time** epochs whose seeds the attacker cannot replay. Thus, **quantum cannot help**: it neither introduces bias into the data nor obtains an oracle with sufficient bandwidth to matter.

---

15. Operational Requirements to Preserve Perfect Secrecy

Shannon perfect secrecy is a **mathematical** condition, but its preservation in practice depends on operational discipline:

1. **Fresh one-time projection P per epoch.** Do not reuse P across distant sessions; cross-epoch reuse can, in theory, allow correlation if other controls fail.
2. **Uniform, independent masks Z.** Any bias or predictability in Z would break the algebra $P(Y = y \mid M = m) = 1/6$.
3. **Balanced leaves and diffusion.** Ensure that the allocation of symbols to leaves is as equal as possible and that ring-rotations are fresh. Minor imbalances are tolerable if they are independent of M; but engineering strives for equal loads.
4. **Isolate epoch identifiers and hash seeds.** Public metadata should commit to freshness without revealing seeds sufficient to invert P or R.
5. **UI/Timing hygiene.** Quantize response windows; add jitter where helpful; suppress side channels.

With these, the perfect-secrecy proofs apply as written.

---

16. Synthesis: Perfect Secrecy by Design, Not by Difficulty

It is tempting to conflate **hardness** with **secrecy**. This dissertation stresses the opposite: **secrecy stems from independence**, engineered by one-time randomized projection P and per-round masking Z; **difficulty** is merely a fallback when someone insists on enumerating states anyway.

- **Secrecy claim:** $I(M;C) = 0$ because $Y = S \oplus Z$ with $S = P(M)$ uniform and Z uniform.
- **No residue:** all observed values are equally probable; there is no frequency or correlation structure to mine.
- **Brute-force crater:** private map 6!, ring rotations $x = 10800$ per applicable degree, yielding $x^{6!}$ in the upper-bound model or $720 \cdot x^h$ in the per-round model, both enormous.
- **Quantum impotence:** without structure or an oracle, no advantage exists; even with Grover fantasy, exponents remain prohibitive.

The human/agent **efficiency** sits orthogonally to these guarantees: perceptual Gestalt search across six leaves and constant-time masking operations make proving **fast and natural**. The attacker's view remains white noise.

---

17. Conclusion

We have shown, with formal statements and practical elaboration, how the combination of a **one-time, entropy-dependent projection** (OTP), **ring-rotation diffusion** across an $x = 30 \times 30 \times 12$ configuration space, **confusion** via a **private six-way bijection** and its 6! **second-layer** permutations, and **masked responses** yields **Shannon-style perfect secrecy** for the interactive transcript. The equality

$$P(M = m \mid C = c) = P(M = m) \quad \forall m, c$$

"The probability of the message given the transcript equals the prior probability of the message for every message and transcript."

holds because each observed response is, by construction, **uniform** over six possibilities and **independent** of the hidden symbol. Consequently, the attacker obtains **no statistical residue**, no biases, no correlations, no frequencies, from which to infer anything. If the attacker fantasizes about guessing the

hidden state instead, they face spaces that scale at least as $720 \cdot 10800^h$ and, in the natural upper bound, as $10800^{720} \approx 10^{2904}$. The quadratic speedup of Grover's algorithm leaves these emphatically infeasible; and more importantly, **quantum computation cannot informatively post-process statistically independent data**.

The system's utility for honest provers, human or agent, derives from a clean asymmetry: the prover holds the private map and knows the next symbol, enabling **Gestalt-fast** localization and **instant** masked responses; the adversary holds neither and sees only featureless noise. This is the heart of "knowing without showing": the *right* knowledge makes the proof effortless; the *wrong* vantage reduces the world to uniformity.

In sum, the design achieves what Shannon prescribed: **all values are equally probable** at the observable interface. The attacker gains nothing, computationally, statistically, or quantum-mechanically, because there is nothing to gain.