ENI6MA Witness Complexity

ENI6MA WITNESS COMPLEXITY

1 · Set-up and quick numbers

Let

Symbol	Meaning	Typical value in your scenario
α	alphabet size in one challenge state	$\alpha = 4$ (the "1-in-4" pick)
ν	witness length inside one interactive proof	$\nu = 4$ states
ρ	how many *complete* ν-state transcripts you chain	$\rho = 4$ repetitions

2 · Permutation space and success probability

A single transcript exposes one ν-symbol word over an alphabet of size α. Combinatorically

$$\Omega(\alpha, \nu) \;=\; \alpha^{\,\nu}, \qquad P_{\text{one}}(\alpha, \nu) \;=\; \alpha^{-\nu}.$$

Chaining ρ independent transcripts multiplies exponents:

$$\Omega_{\text{total}} = \alpha^{\,\nu\rho}, \qquad P_{\text{total}} = \alpha^{-\nu\rho}$$

Plugging the concrete values

$$\alpha^{\nu\rho} = 4^{4\times4} = 4^{16} = 4{,}294{,}967{,}296, \qquad P_{\text{total}} = 4^{-16} \approx 2.3 \times 10^{-10}.$$

So every **online guess succeeds only once in ~4 billion trials**; an offline attacker must enumerate four billion projected-witness candidates *per genuine password*.

---

3 · Why this "simple" exponential matters in practice

1. **Linear cost per candidate → exponential wall** The attacker's inner loop—evaluate the verifier's check for a guessed witness—still costs essentially the same whether you show one transcript or four. But concatenating transcripts multiplies the *count* of guesses by $\alpha^{\nu(\rho-1)}$. Four repeats lift effort from $4^4 = 256$ to $4^{16} \approx 4.3$ billion—already beyond what real-time online throttling or even modest rate limits will allow.

2. **Lower-dimensional leakage is not exploitable** Each state is a **projection bit** (membership of the secret vector S in a private hyper-plane). Even though the witness lives in a 1-bit manifold, the alphabet that labels *which* plane is queried varies round-to-round on a modulo-indexed ring of alphabets $\Sigma_i = \{\sigma_{i,0}, \ldots, \sigma_{i,\alpha-1}\}$. Because the ring evolves quasi-randomly, the attacker never sees two states whose indices are correlated in a way that would let them harvest frequencies or phase information. Formally, for indices $j \neq k$

$$\Pr[\sigma_{i,j} = \sigma_{i+\delta,k}] = \frac{1}{\alpha},$$

so no fixed offset δ yields an exploitable bias.

3. **No shortcut through "clever" cross-round inference** Each projection bit is XOR-mixed with a fresh mnemonic mask in the prover's head **before** transmission. The mutual information

$$I(S;\, w_i) = 0,$$

meaning that—even after four rounds—the attacker learns nothing about the password other than "it exists".

---

4 · Intuition: stacking thin slices into an iron wall

Picture the secret password as a point hidden inside a 4-D crystal. *One* witness slice (one transcript) is a sheet of colored glass: peering through it, the adversary sees only one colour among four. Stacking four sheets at slightly rotated angles multiplies the patterns a would-be code-breaker must match: every extra sheet quadruples the palette, and after four sheets the spectrum has 4^4 distinct hues. Ask the prover to *repeat* the entire stack four times, and the attacker faces 4^{16} colour codes. Trying them all is no longer guesswork—it is geological time.

---

5 · Take-away

The "repeat-the-PoK" tweak looks trivial, yet mathematically it elevates the brute-force search from hundreds to **billions** of combinations while preserving the zero-knowledge veil that makes each individual bit worthless in isolation. The modulo-ring alphabet schedule blocks statistical footholds, so the lower-dimensional witness leaks no structure the attacker could chain across repetitions. In short, four 1-in-4 challenges, asked four times, turn a toy-sized attack surface into a mountain no practical adversary can climb.

Attack Complexity On ZK Witness via Interactive Sigma Protocol Proofs

1 · Set-up and quick numbers

Let

Symbol	Meaning	Typical value in your scenario
α	alphabet size in one challenge state	$\alpha = 4$ (the "1-in-4" pick)
ν	witness length inside one interactive proof	$\nu = 4$ states
ρ	how many *complete* ν-state transcripts you chain	$\rho = 4$ repetitions

2 · Permutation space and success probability

A single transcript exposes one ν-symbol word over an alphabet of size α. Combinatorically

$$\Omega(\alpha, \nu) \;=\; \alpha^{\,\nu}, \qquad P_{\text{one}}(\alpha, \nu) \;=\; \alpha^{-\nu}.$$

Chaining ρ independent transcripts multiplies exponents:

$$\Omega_{\text{total}} = \alpha^{\,\nu\rho}, \qquad P_{\text{total}} = \alpha^{-\nu\rho}$$

Plugging the concrete values

$$\alpha^{\nu\rho} = 4^{4\times4} = 4^{16} = 4{,}294{,}967{,}296, \qquad P_{\text{total}} = 4^{-16} \approx 2.3 \times 10^{-10}.$$

So every **online guess succeeds only once in ~4 billion trials**; an offline attacker must enumerate four billion projected-witness candidates *per genuine password*.

3 · Why this "simple" exponential matters in practice

1. **Linear cost per candidate → exponential wall** The attacker's inner loop—evaluate the verifier's check for a guessed witness—still costs essentially the same whether you show one transcript or four. But concatenating transcripts multiplies the *count* of guesses by $\alpha^{\nu(\rho-1)}$. Four repeats lift effort from $4^4 = 256$ to $4^{16} \approx 4.3$ billion—already beyond what real-time online throttling or even modest rate limits will allow.

2. **Lower-dimensional leakage is not exploitable** Each state is a **projection bit** (membership of the secret vector S in a private hyper-plane). Even though the witness lives in a 1-bit manifold, the alphabet that labels *which* plane is queried varies round-to-round on a modulo-indexed ring of alphabets $\Sigma_i = \{\sigma_{i,0}, \ldots, \sigma_{i,\alpha-1}\}$. Because the ring evolves quasi-randomly, the attacker never sees two states whose indices are correlated in a way that would let them harvest frequencies or phase information. Formally, for indices $j \neq k$

$$\Pr[\sigma_{i,j} = \sigma_{i+\delta,k}] = \frac{1}{\alpha},$$

so no fixed offset δ yields an exploitable bias.

3. **No shortcut through "clever" cross-round inference** Each projection bit is XOR-mixed with a fresh mnemonic mask in the prover's head **before** transmission. The mutual information

$$I(S;\, w_i) = 0,$$

meaning that—even after four rounds—the attacker learns nothing about the password other than "it exists".

---

4 · Intuition: stacking thin slices into an iron wall

Picture the secret password as a point hidden inside a 4-D crystal. *One* witness slice (one transcript) is a sheet of colored glass: peering through it, the adversary sees only one colour among four. Stacking four sheets at slightly rotated angles multiplies the patterns a would-be code-breaker must match: every extra sheet quadruples the palette, and after four sheets the spectrum has 4^4 distinct hues. Ask the prover to *repeat* the entire stack four times, and the attacker faces 4^{16} colour codes. Trying them all is no longer guesswork—it is geological time.

---

5 · Take-away

The "repeat-the-PoK" tweak looks trivial, yet mathematically it elevates the brute-force search from hundreds to **billions** of combinations while preserving the zero-knowledge veil that makes each individual bit worthless in isolation. The modulo-ring alphabet schedule blocks statistical footholds, so the lower-dimensional witness leaks no structure the attacker could chain across repetitions. In short, four 1-in-4 challenges, asked four times, turn a toy-sized attack surface into a mountain no practical adversary can climb.

Consolidated Mathematical Expression for Brute-Force Probability:

Consider a scenario involving multiple rounds n, each having a discrete set of hyperplanes with size s. If an attacker tries to brute force the witness directly, the total probability P of guessing the correct witness across all rounds is given by the general expression:

$$P = \prod_{i=1}^{n} \frac{1}{s_i}$$

Where:

- n is the number of rounds,
- s_i is the number of distinct hyperplanes (or states) possible in the i^{th} round.

For a simplified uniform scenario, where every round has the same number of distinct hyperplanes s, this reduces to:

$$P = \left(\frac{1}{s}\right)^n$$

For instance, with a set of $s = 4$ hyperplanes across $n = 6$ rounds, the brute-force probability is explicitly computed as:

$$P = \left(\frac{1}{4}\right)^6 = \frac{1}{4096}$$

Explanation and Limitations of the Naive Brute-Force Approach:

While the above permutation-based calculation is straightforward and intuitive, it significantly underestimates the complexity and inherent protection provided by the cryptographic mechanism in question.

1. Stochastic Alphabet Distributions and Entropy:

The hyperplane selections (witness memberships) are not purely deterministic or uniformly distributed. Instead, they follow a stochastic distribution determined by the **entropy** of the **manifold projection**, where entropy represents the measure of uncertainty and randomness inherent to the system. The actual probability of correctly guessing is not equally likely for every hyperplane because the mappings between hyperplanes and the alphabet distributions are influenced by highly unpredictable transformations.

2. Quantum Entropy and Manifold State Collapse:

Prior to verification, the state of the system—particularly within holographic projections—is not fully determined. The system can exist simultaneously in multiple potential states (analogous to quantum superposition). Upon measurement or verification (the “collapse”), one of these states is selected. This quantum entropy and uncertainty mean that even the theoretical knowledge of permutations does not translate into practical advantage. Until verification (measurement), the manifold projection is inherently uncertain, and this uncertainty is irreducible through classical brute-force attempts.

Mathematically, the entropy-based uncertainty H within the quantum manifold projection can be expressed as:

$$H = -\sum_{x \in X} p(x) \log p(x)$$

Here, X is the alphabet set representing all potential hyperplane states, and $p(x)$ is the probability distribution over these states. Due to the highly stochastic and uncertain nature of $p(x)$, brute force provides no meaningful advantage—each guess effectively remains random.

3. Opaque Mapping and XOR Projection:

The mnemonic map (private map) used for verification operates through an XOR calculation embedded cognitively in the authenticator's mind. This XOR-based mapping obscures the relationship between the hyperplanes and the secret witness sequence. Without knowledge of the secret XOR keys (the authenticator's mental map), brute forcing the hyperplane states alone cannot derive the underlying secret password or witness.

4. Implication for Security:

Because each guess of the witness membership is effectively independent and equally uncertain without direct knowledge of the mnemonic map, the actual security is exponentially enhanced beyond a simple permutation analysis. In other words, even a computationally exhaustive attack on hyperplane memberships merely enumerates possibilities without any incremental knowledge about the correct mnemonic sequence. This leaves an attacker no closer to identifying the secret password or witness sequence after brute-force efforts than before.

Why the Simple Permutation Argument is Misleading:

In summary, the simple combinational or permutation-based brute-force calculation severely oversimplifies the scenario. It ignores:

- **Non-uniform stochastic distributions** due to the inherent entropy and randomness.
- **Quantum uncertainty and indeterminate states** that exist prior to manifold state collapse.
- **Opaque XOR mappings** that prevent incremental discovery of the secret through brute-force guessing alone.

Consequently, the actual security against a brute-force attack is significantly stronger than indicated by basic combinational mathematics. The complexity introduced by these advanced cryptographic principles ensures that brute-force guessing provides no practical advantage in discovering the witness or the mnemonic map sequence.

1 · Consolidated probability expression

Let each authentication round draw a **random hyper-plane** $H_i \subset \mathbb{R}^d$ and a **stochastic alphabet** $\Sigma_i = \{\sigma_{i,1}, \dots, \sigma_{i,K_i}\}$ whose size K_i is itself a random variable driven by a private entropy source. The secret password is a latent vector $S \in \mathbb{R}^d$. The public verifier sees only the **witness**

$$w_i \;=\; \phi_i(S, \xi_i) \;=\; \begin{cases} 1 & \text{if } S \in H_i \\ 0 & \text{otherwise} \end{cases}$$

where ξ_i is fresh, unshared entropy. Because ϕ_i hides both H_i and ξ_i, the attacker's best strategy in round i is a blind guess among the K_i equally probable symbols.

Denote the full witness sequence by $\mathbf{w} = (w_1, \dots, w_n)$. Conditioned on the private map, the rounds are independent, so the success probability of a brute-force attack that tries to guess the entire witness is

$$P_{\text{attack}} = \Pr\big[\hat{\mathbf{w}} = \mathbf{w}\big] = \prod_{i=1}^{n} \frac{1}{K_i} = \exp\Big(-\sum_{i=1}^{n} \ln K_i\Big). \tag{1}$$

If every round happened to use the same alphabet size K, (1) collapses to the familiar K^{-n}.

---

2 · Why the naïve "K^{-n}" intuition is misleading

Naïve assumption	Reality in the ENI6MA manifold projection
K_i is fixed and public	K_i is *random* and hidden; its distribution comes from quantum-grade entropy sampled inside the circuit authority.
Hyper-planes are static	Each H_i is re-drawn per round from a high-entropy measure on $\mathbb{R}^d$, so there is **no correlation** an attacker can exploit across rounds.
The witness reveals location information	The witness is a single bit per round and is XOR-combined in the user's mind before transmission, leaking *zero mutual information* about S.
Search space is exactly $\prod K_i$	Even if the attacker enumerated all $\prod K_i$ candidates, they would still not know which (if any) corresponds to a valid password, because the mapping ϕ_i is private and one-way. Their work factor grows only **linearly** (by an "N" multiplier) beyond conventional guessing but yields *no additional knowledge*.

---

3 · Intuition: the hyper-plane as a one-bit hologram

Think of S as a light source and each H_i as a slide in a projector. The verifier sees only the on/off sparkle that leaks through the slide, **not** the slide's pattern or the light-source coordinates. Replacing the slides every round—with layouts chosen by quantum entropy—turns the attacker's view into a sequence of coin flips whose bias is unknown and irreproducible.

---

4 · Consequence for brute-force strategy

1. **Indistinguishability** All guesses that pass the public check are computationally indistinguishable from random strings; no ranking is possible.
2. **Linear blow-up, zero gain** Brute-forcing the witness merely multiplies the attacker's work by $N = \mathbb{E}[\prod K_i]$ but does **not** reduce the uncertainty about the hidden password.
3. **Entropy preservation** Because the XOR aggregation happens *before* any witness leaves the user's cognition, the total Shannon entropy exposed to the network is exactly one bit per round—irreducible under any classical or quantum analysis.

---

5 · Bottom line

Equation (1) is already pessimistic for the defender; in practice the attacker does not even know the K_i. The manifold's stochastic projection plus the private map converts every round into an information-theoretic blind trial. Consequently, targeting the witness sequence rather than the password yields **no shortcut:** the expected success probability collapses to the same exponential bound, and every additional round drives it closer to zero at a rate determined by fresh quantum entropy rather than by any property the attacker can observe or predict.