Computational and Quantum Intractability for ENI6MA Membership-Only Proofs of Knowledge

by Frank Dylan Rosario and Dr. Lin Wang

Abstract

We present an analysis of a membership-only proof-of-knowledge protocol whose observable transcript is limited to six-way leaf identifiers. The design uses balanced partitions of a 72-symbol alphabet, independently rotating alphabet rings, and a private bijection over six labels to throttle information leakage per round while deliberately stripping away exploitable structure. We formalize the attacker's problem as identification of a marked hypothesis in a combinatorial space of size $|\mathcal{H}| = 6!\binom{U}{6} = P(U, 6)$, with $U = |\Sigma|^L = 72^L$ for secret length L. For six distinct secrets of length $L = 6$, the hypothesis count is $|\mathcal{H}| \approx 72^{36} \approx 2^{222}$; for $L = 12$, $|\mathcal{H}| \approx 72^{72} \approx 2^{444}$. We prove information-theoretic lower bounds on the number of rounds required to isolate a unique solution ($R_{\min} \approx \lceil \log_6 |\mathcal{H}| \rceil$: about 86 rounds when $L = 6$ and about 172 when $L = 12$), and we give black-box lower bounds on classical time $\Omega(R\,|\mathcal{H}|)$ and quantum time $\Omega(R\sqrt{|\mathcal{H}|})$ via BBBV/Grover limits. We then convert these asymptotics into resource models that explicitly account for reversible oracle construction, fault-tolerant overheads, and physical-layer constraints. The conclusion is robust: even with fantastically optimistic hardware assumptions, the protocol remains computationally intractable for both $L = 6$ and $L = 12$, with the $L = 12$ setting squaring the already astronomical work factor.

---

1. Introduction

Contemporary authentication systems wrestle with a tradeoff between secrecy and verifiability. Traditional passwords, biometrics, and long-lived keys leak enough structure that adversaries can accumulate, correlate, and replay fragments over time. Algebraic zero-knowledge (ZK) protocols resolve this by furnishing verifiable statements without revealing witnesses, but often at the cost of sophisticated assumptions, heavy arithmetic, and complex implementation.

The protocol analyzed here occupies an orthogonal point in the design space: it deliberately minimizes **semantic output**—reducing each round to a single six-way label—while engineering the surrounding geometry (balanced partitions, independent rotations, and a private six-way bijection) so that an observer's transcript is, in expectation, **statistically indistinguishable from uniform noise**. The result is that no gradient, bias, or algebraic scaffold survives for an attacker to climb; all that remains is eliminative consistency checking over an enormous, structureless hypothesis set.

Two concrete parameterizations ground the discussion. In both, Alice (the prover) possesses **six distinct secrets**. In the baseline model, each secret has length $L = 6$; in the expanded model, $L = 12$. At each round, a public board partitions the 72-symbol alphabet into six equal "leaves," while the three alphabet rings (lowercase, uppercase, digits) undergo independent modulo rotations. Alice's response is a codeword drawn from a private six-way bijection that permutes the leaf labels. To a passive eavesdropper, the visible label per round has a uniform marginal distribution on $\{1, \ldots, 6\}$, and across rounds the independent rotations sever positional anchors. Across ceremonies, independent choices of private bijections preserve **label-switching symmetry**. In short, everything an attacker might use to carve down the search space is intentionally washed out; we are left with **unstructured search**.

This dissertation pursues three goals. First, it formalizes the attacker's task as marked-item identification in a combinatorial hypothesis set and gives exact counts for the size of that set at $L = 6$ and $L = 12$. Second, it quantifies the information content per round and derives the rounds-to-uniqueness thresholds that any passive observer must meet to isolate a unique global solution. Third, it translates black-box classical and quantum lower bounds into resource models that include the real costs of building **reversible membership oracles** and operating them under **fault-tolerant quantum error correction**. The main message is consistent across all three: the construction is **computationally and quantumly intractable** at the baseline, and the expanded length doubles the exponent and thereby squares the work.

---

2. Entities, Alphabet Geometry, and Observable Transcript

Let the alphabet be $\Sigma = \Sigma_1 \dot{\cup} \Sigma_2 \dot{\cup} \Sigma_3$ with $(|\Sigma_1|, |\Sigma_2|, |\Sigma_3|) = (30, 30, 12)$ and total $|\Sigma| = 72$. A secret is a word $C \in \Sigma^L$ of length L. Alice holds six distinct secrets $S = \{C^{(1)}, \ldots, C^{(6)}\} \subset \Sigma^L$. Each round i is parameterized by public randomness B_i: three independent ring rotations $(\Delta_i^{(1)}, \Delta_i^{(2)}, \Delta_i^{(3)})$ and a balanced partition $\Pi_i : \Sigma \to \{1, \ldots, 6\}$ that places exactly 12 symbols in each leaf. The ring rotations send in-ring indices $j \mapsto j' = (j + \Delta_i^{(r)}) \bmod |\Sigma_r|$, independently per ring.

The next character Alice must demonstrate lies in some secret $C^{(k)}$ at posi-

tion t; its current leaf is $L_i := \Pi_i(C_t^{(k)})$. Alice's observable reply is a codeword $Y_i := \varphi^{-1}(L_i)$, where $\varphi \in S_6$ is a **private bijection** pairing her six codewords with the six leaves. To the outside observer, who lacks φ, the single symbol Y_i is uniformly distributed on six outcomes once B_i is fixed; averaged over φ, the distribution of Y_i is uniform regardless of $C_t^{(k)}$. Across rounds, fresh rotations of each ring erase positional correlation. Across ceremonies, fresh or opaque bijections erase label identities. The transcript available to an attacker is therefore a sequence $\{(\mathsf{B}_i, Y_i)\}_i$ in the board-visible model, or just $\{Y_i\}_i$ in a board-hidden variant. In either case, the transcript's **marginals are uniform** and the **mutual information about the next character is, in expectation, zero**.

---

3. Hypothesis Space and Exact Counting

Write $U = |\Sigma|^L = 72^L$ for the number of possible secrets of length L. Because Alice's secrets are distinct and because φ can be any of the 6! permutations, the global hypothesis set is

$$\mathcal{H} = \{(S, \varphi) : S \subset \Sigma^L,\ |S| = 6,\ \varphi \in S_6\}.$$

Counting is exact and instructive:

$$|\mathcal{H}| = 6! \binom{U}{6} = P(U, 6) = U(U-1)(U-2)(U-3)(U-4)(U-5).$$

This identity says that "unordered six secrets + a six-way bijection" is equinumerous with "ordered six distinct secrets"; the private map simply provides an order. When $U \gg 1$, Stirling's approximation gives $P(U, 6) \sim U^6$. Passing to base-2 logarithms,

$$\log_2 |\mathcal{H}| \approx 6 \log_2 U = 6L \log_2 72 \approx 6L \times 6.17 \approx 37.0\, L \ \text{ bits}.$$

Two concrete instantiations highlight the growth:

- **Six-character secrets ($L = 6$):** $|\mathcal{H}| \approx 72^{36} \approx 2^{222}$.
- **Twelve-character secrets ($L = 12$):** $|\mathcal{H}| \approx 72^{72} \approx 2^{444}$.

Doubling the secret length doubles $\log_2 U$ and thus doubles $\log_2 |\mathcal{H}|$; equivalently, it **squares** $|\mathcal{H}|$. The move from $L = 6$ to $L = 12$ therefore multiplies the attacker's candidate set by roughly 2^{222}—a second "astronomical" factor atop the first.

---

4. Information-Theoretic Leakage and Rounds-to-Uniqueness

Because each partition is balanced and each ring is independently re-indexed every round, the unconditional marginal over observed labels is flat. Formally, for any fixed board B_i and any symbol $c \in \Sigma$,

$$\Pr(Y_i = j \mid C_i = c, \mathsf{B}_i) = \tfrac{1}{6} \quad \text{for all } j \in \{1, \ldots, 6\},$$

so the per-round mutual information is bounded by the entropy of a six-way outcome,

$$I(C_i; Y_i \mid \mathsf{B}_i) \le H(Y_i \mid \mathsf{B}_i) = \log_2 6 \approx 2.585 \text{ bits}.$$

Averaging over φ and the independent rotations, the **expected** mutual information about the *next* character is effectively zero; there is no learnable directional signal. In a single continuous session with a fixed but hidden φ, however, the transcript **does** allow eliminative consistency checking: wrong hypotheses remain compatible with the next round with probability $\approx 1/6$. After R independent rounds, the expected number of wrong survivors is $|\mathcal{H}|(1/6)^R$, and "uniqueness in expectation" requires

$$|\mathcal{H}|(1/6)^R \lesssim 1 \quad \Longleftrightarrow \quad R \gtrsim \log_6 |\mathcal{H}| = \frac{\log_2 |\mathcal{H}|}{\log_2 6}.$$

This yields the concrete thresholds:

$$R_{\min}(L = 6) \approx \left\lceil \frac{222.1}{2.585} \right\rceil \approx 86, \qquad R_{\min}(L = 12) \approx \left\lceil \frac{444.2}{2.585} \right\rceil \approx 172.$$

These are **information-theoretic** lower bounds that assume perfect visibility of every round in one uninterrupted session. Any practical obscurity (hidden boards, subsampling, decoy rounds) increases the required effort for the attacker or allows the defender to stop earlier for the same safety.

---

5. Classical Black-Box Complexity

If the only operation available to the attacker is to test a hypothesis (S, φ) against a transcript, then the problem is a **black-box** exhaustive search over $|\mathcal{H}|$ items using a membership test of cost $O(R)$. By decision-tree lower bounds (and Yao's minimax principle for randomized algorithms), the expected number of tests needed to find the unique marked item with constant success probability is $\Omega(|\mathcal{H}|)$. Therefore the classical time is

$$T_{\text{classical}}(L, R) = \Omega\big(R\,|\mathcal{H}|\big) = \Omega\big(R\,72^{6L}\big).$$

At $L = 6$, this is $\Omega(R\,72^{36}) = \Omega(R\,2^{222})$; at $L = 12$, it is $\Omega(R\,72^{72}) = \Omega(R\,2^{444})$. Interpreting these exponents as work factors makes the intractability evident. Even granting a physically implausible 10^{15} consistency checks per second and ignoring memory and I/O, scanning 2^{222} candidates would require $\sim 1.9 \times 10^{44}$ years; squaring that space for $L = 12$ is simply beyond human metaphor. In practice the per-candidate check includes parsing and mapping across R rounds; thus constants are unfavorable as well.

---

6. Quantum Query Lower Bounds and Grover-Regime Costing

Quantum algorithms cannot asymptotically beat unstructured exhaustive search except by a square root. The BBBV lower bound shows that any quantum algorithm that identifies a marked item in a set of size M with bounded error must make $\Omega(\sqrt{M})$ queries to an oracle that recognizes the marked item. In our setting, $M = |\mathcal{H}| = P(U, 6)$, and the appropriate oracle takes a candidate (S, φ) and produces a predicate indicating consistency with the transcript. Querying this oracle once costs $\Theta(R)$ reversible steps (more in a fault-tolerant setting) because it must compute each round's predicted label and compare it to the observed one inside a reversible circuit, leaving the workspace clean.

Consequently, quantum time satisfies

$$T_{\text{quantum}}(L, R) = \Omega\big(R\,\sqrt{|\mathcal{H}|}\big) = \Omega\big(R\,72^{3L}\big).$$

The square-root improvement cuts the exponent in half but leaves the numbers astronomical. For $L = 6$, $\sqrt{|\mathcal{H}|} \approx \sqrt{72^{36}} = 72^{18} \approx 2^{111}$; for $L = 12$, $\sqrt{|\mathcal{H}|} = 72^{36} \approx 2^{222}$. Suppose, wildly optimistically, that a fully error-corrected quantum computer could execute 10^{18} reversible oracle calls per second with negligible constant factors. Then 2^{111} oracle calls would still require $\sim 7.9 \times 10^7$ years; and that is before accounting for the substantial overhead of synthesizing the oracle and running it fault-tolerantly. Parallelizing Grover on p independent quantum processors yields at most a $\sqrt{p}$ speedup; even $p = 10^{12}$ (a trillion) only buys a factor of 10^6, leaving centuries to millennia under implausibly perfect conditions.

The critical point is that the **oracle itself** is not free: it must compute, within a reversible circuit, the leaf membership for each round given candidate secrets and candidate φ, compare the results to the recorded labels, and restore ancillas to zero. That introduces depth and width overhead proportional to R, which grows linearly with L. Therefore, as L moves from 6 to 12, quantum tractability degrades in two ways: the **exponent doubles** (from 2^{111} to 2^{222} oracle calls), and the **per-oracle cost increases** with the longer transcript.

---

7. Fault-Tolerant Quantum Cost Model (T-Count, T-Depth, and Logical Qubits)

A more realistic quantum estimate assigns costs to a reversible membership oracle in a surface-code regime. One encodes logical qubits at a target physical error rate and compiles the oracle to Clifford+T gates, accounting for T-count and T-depth, Toffoli synthesis, and ancilla management. The oracle's structure is straightforward: for each round i, compute (reversibly) the predicted label $\tilde{Y}_i$ from the candidate (S, φ) and the public B_i, then apply a phase-flip conditioned on $\tilde{Y}_i = Y_i$. Each such subroutine is dominated by address computation, modular addition for ring rotations, table-like partition lookups (which can be hashed or index-computed), and reversible comparison. Because the board partitions are balanced but arbitrary, one cannot exploit fixed arithmetic structure; at best, one uses small reversible lookup gadgets or arithmetic that simulates the balanced mapping.

A coarse scaling law is enough for our purposes: if a single round compiles to $O(1)$ Toffoli-equivalents, then an R-round oracle compiles to $\Theta(R)$ Toffolis, with T-count and T-depth linear in R, and logical-qubit footprint growing with the number of simultaneously evaluated rounds or with the degree of pebbling used to trade space for depth. Since $R \approx 86$ at $L = 6$ and $R \approx 172$ at $L = 12$, the reversible oracle's resource usage roughly doubles when doubling L. In Grover's algorithm, the total T-count is the per-oracle T-count times the number of iterations, and the total runtime is the per-oracle time times the iteration count. Thus, at $L = 12$ the **iteration count doubles in the exponent** and the **per-iteration cost roughly doubles** in the linear factor. The product is devastating for feasibility.

---

8. Physical-Layer Sanity Checks: Time, Energy, and I/O

It is instructive to complement asymptotic bounds with crude but telling physical estimates. Imagine a classical engine that could test 10^{15} hypotheses per second—already beyond what you can sustain for nontrivial membership tests with realistic memory hierarchies. The wall-clock time for $L = 6$ is $\sim 2^{222}/10^{15}$ seconds, or on the order of 10^{44} years; for $L = 12$ it is $\sim 2^{444}/10^{15}$ seconds, a time dwarfing astrophysical scales. Energy considerations via Landauer's bound only worsen the picture: even if each test cost the erasure of a single bit (it does not), the energy to touch 2^{222} hypotheses would be unimaginably large. I/O and memory pressure dominate long before arithmetic does: enumerating or streaming candidate sextuples and private maps at these magnitudes is itself a bottleneck. On the quantum side, magic-state distillation for T-gates and the need to maintain large numbers of logical qubits coherently for years to millennia render optimistic Grover-regime estimates fanciful for both $L = 6$

and $L = 12$. The point of these checks is not precision; it is to ground the lower-bound exponents in engineering reality.

---

9. The Board-Hidden Variant and Multiplicative Explosion

Thus far we have assumed a board-visible model in which the attacker sees B_i. If, instead, boards are hidden (e.g., via a secure attention window), then the attacker must **also** hypothesize the per-round ring rotations and, if not derivable, the balanced partition itself. The three rings admit $30 \times 30 \times 12 = 10{,}800$ independent rotation triples per round, multiplying the hypothesis space by $10{,}800^R$. If the balanced partition is not public deterministically (e.g., if it is derived from hidden randomness), then per round the attacker faces a multinomial count

$$\frac{72!}{(12!)^6\, 6!}$$

for leaf assignments, another astronomical multiplicative factor. Either way, hiding or salting board details only **increases** the search burden relative to the already intractable board-visible baseline.

---

10. Why Statistics, Correlations, and Learning Fail

Cryptanalysis often thrives on structure: bias in S-boxes, linear or differential trails, algebraic relations, or repeated keystreams. This design neutralizes such advantages. Balanced partitions force leaf marginals to 1/6 per round; independent ring rotations obliterate cross-round positional anchors; and the private bijection makes labels exchangeable, so that a transcript is consistent with 6! relabelings. Across ceremonies, frequencies converge to uniform and label identities do not carry across sessions. In the limit, the best any learning procedure can do on passive transcripts is to model a six-sided die. Intra-session eliminative consistency is the only viable path—and that reduces to the black-box search already analyzed. The absence of exploitable signal is not accidental; it is engineered.

---

11. Comparative Models: $L = 6$ vs. $L = 12$

It is now straightforward to juxtapose the two length regimes along key axes:

Hypothesis size. The count $|\mathcal{H}| = P(U, 6)$ satisfies $\log_2 |\mathcal{H}| \approx 6L \log_2 72 \approx 37.0L$. Moving from $L = 6$ to $L = 12$ doubles $\log_2 |\mathcal{H}|$ from ~ 222 to ~ 444; equivalently, $|\mathcal{H}|$ is **squared**.

Rounds-to-uniqueness. The information bound yields $R_{\min} \approx \lceil \log_6 |\mathcal{H}| \rceil$, thus roughly doubling from ~ 86 to ~ 172 as L doubles, because per-round leakage is capped by $\log_2 6$ bits.

Classical time. $\Omega(R\,|\mathcal{H}|)$ scales as $\Omega(R\,72^{6L})$, so the move to $L = 12$ multiplies time by about $2 \cdot 72^{36} \approx 2 \cdot 2^{222}$—an extra factor of roughly 2^{223} over the already infeasible baseline.

Quantum time. $\Omega(R\sqrt{|\mathcal{H}|})$ scales as $\Omega(R\,72^{3L})$; moving to $L = 12$ multiplies time by about $2 \cdot 72^{18} \approx 2^{112}$. Moreover, the per-oracle reversible cost is roughly doubled in R.

Fault-tolerant overhead. The number of rounds R enters linearly in the oracle's T-count, T-depth, and ancilla footprint. The iteration count for Grover climbs from $\Theta(2^{111})$ to $\Theta(2^{222})$. The product of these two growths—linear in R, exponential in L—pushes any credible resource estimate well beyond technological horizons.

The gestalt is clear: the $L = 12$ regime is not simply "more secure;" it is a **regime change** in which the already astronomical search space is squared, the rounds double, and every component of the quantum resource model worsens accordingly.

---

12. Parameter Tuning and Security Targets

Security discussions often target "bits of work." For classical adversaries, **128-bit security** means 2^{128} steps. At $L = 6$, $\log_2 |\mathcal{H}| \approx 222$ already exceeds this by ~ 94 bits. For Grover-class quantum adversaries, one often seeks **post-quantum 128-bit** security, i.e., $\sqrt{|\mathcal{H}|} \geq 2^{128}$, equivalently $|\mathcal{H}| \geq 2^{256}$. The $L = 6$ setting falls short of this formal bar (though it remains astronomically hard in practice, as above); the smallest increase $L = 7$ yields $|\mathcal{H}| \approx 72^{42} \approx 2^{259}$ and $\sqrt{|\mathcal{H}|} \approx 2^{129.5}$, clearing the bar cleanly. The $L = 12$ regime far exceeds it: $\sqrt{|\mathcal{H}|} \approx 2^{222}$. Designers have additional knobs beyond length—number of secrets and alphabet size—and can trade usability for security by adjusting any of them. The formulas are compact:

$$|\mathcal{H}| \approx |\Sigma|^{6L}, \quad \log_2 |\mathcal{H}| \approx 6L \log_2 |\Sigma|, \quad R_{\min} \approx \left\lceil \frac{\log_2 |\mathcal{H}|}{\log_2 6} \right\rceil.$$

These encapsulate the architecture's tunability.

---

13. Robustness, Edge Cases, and Threat Extensions

If users weaken the entropy source—e.g., by selecting secrets from a tiny dictionary $D \ll U$—then $|\mathcal{H}|$ falls to $P(D, 6)$ and the protection collapses to the dictionary's entropy. That is a policy issue, not a structural flaw. Allowing repetitions among the six secrets replaces $P(U, 6)$ with U^6, which is asymptotically equivalent for our purposes. Partial leakage of φ reduces a constant factor 6! to something smaller; this does not materially change $|\mathcal{H}| \approx U^6$. Active attackers (man-in-the-middle) invite standard countermeasures (authenticated channels, nonces, transcript binding) but do not change the passive intractability analysis. Side-channel defenses (constant-time implementations, masked computations, and noise) preserve the intent that the **semantic** output channel remains the lone observable, uniformly distributed symbol per round. Each of these considerations stresses that the cryptanalytic difficulty is baked into the **combinatorics + throttled information + symmetry** triad; implementation only needs to preserve those properties.

---

14. Formal Statements

For reference, the main claims can be compactly stated.

1. **Exact hypothesis size.** With $U = 72^L$,

$$|\mathcal{H}| = 6!\binom{U}{6} = P(U, 6) = U(U-1)(U-2)(U-3)(U-4)(U-5).$$

2. **Asymptotics.** $|\mathcal{H}| \sim U^6 = 72^{6L} = 2^{6L \log_2 72}$.
3. **Per-round leakage bound.** $I(C_i; Y_i \mid \mathsf{B}_i) \le \log_2 6$ bits and, under the design's randomness, $\mathbb{E}[I(C_i; Y_i \mid \mathsf{B}_i)] \approx 0$.
4. **Rounds to uniqueness.** $R_{\min} \approx \lceil \log_6 |\mathcal{H}| \rceil$, yielding $R_{\min} \approx 86$ at $L = 6$ and $R_{\min} \approx 172$ at $L = 12$.
5. **Classical time.** $T_{\text{classical}} = \Omega(R\,|\mathcal{H}|)$.
6. **Quantum time.** $T_{\text{quantum}} = \Omega(R\sqrt{|\mathcal{H}|})$ (BBBV/Grover).
7. **Board-hidden blowup.** An additional factor of $(30 \cdot 30 \cdot 12)^R = 10{,}800^R$ from rotation uncertainty, and potentially $\frac{72!}{(12!)^6\,6!}$ per round from hidden partitions.

These claims do not rely on algebraic hardness; they are consequences of counting, information bounds, and query lower bounds, all driven by the transcript's deliberately structureless nature.

---

15. Conclusion

The protocol studied reduces each round of a proof-of-knowledge ceremony to a six-way membership label emitted through a private bijection, with the alphabet re-partitioned into balanced leaves under independently rotated rings. This geometry deprives an attacker of statistical drift, algebraic relations, or positional anchors. What remains is eliminative consistency across rounds—pure unstructured search. Exact counting shows that recovering **all six** secrets and the private bijection spans a hypothesis set of size $|\mathcal{H}| = 6!\binom{72^L}{6}$, approximated by 72^{6L}. At $L = 6$, this is $\approx 2^{222}$ with $R_{\min} \approx 86$; at $L = 12$, it is $\approx 2^{444}$ with $R_{\min} \approx 172$. Black-box lower bounds then force classical time $\Omega(R\,|\mathcal{H}|)$ and quantum time $\Omega(R\sqrt{|\mathcal{H}|})$. Converting these exponents into concrete resource models—and incorporating the real costs of reversible oracle construction and fault-tolerant execution—yields timeframes far beyond cosmological scales even under implausibly generous assumptions.

The practical lesson for designers is both simple and powerful. If you ensure (i) balanced partitions per round, (ii) independent ring rotations per round, and (iii) a private bijection over a small, uniform output alphabet, then **information-theoretic throttling** and **combinatorial explosion** conspire to enforce intractability on both classical and quantum attackers. The move from six-character to twelve-character secrets does not merely "add security"; it **squares** an already astronomical search and doubles the sample complexity, while increasing the per-oracle cost in any realistic quantum instantiation. The architecture's virtue is that its guarantees come not from fragile algebraic structure but from counts, symmetry, and leakage caps that are easy to reason about and hard to subvert. That is why, for both $L = 6$ and $L = 12$, and even more so for the latter, brute force—classical or quantum—is the only game in town, and it is a game no attacker can win.

Bibliography of References

1. P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science (SFCS)*, Santa Fe, NM, 1994, pp. 124–134. (users.cs.duke.edu)

2. L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC)*, 1996, pp. 212–219. (arXiv)

3. M. Boyer, G. Brassard, P. Høyer, and A. Tapp, "Tight bounds on quantum searching," *arXiv:quant-ph/9605034*, 1996; published in conference proceedings. (arXiv)

4. C. Zalka, "Grover's quantum searching algorithm is optimal," *Phys. Rev. A*, vol. 60, no. 4, pp. 2746–2751, 1999. (Physical Review)

5. R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf, "Quantum lower bounds by polynomials," in *Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS)*, 1998. (arXiv)

6. C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, "Strengths and weaknesses of quantum computing," *arXiv:quant-ph/9701001*, 1997. (arXiv)

7. A. Montanaro, "Quantum algorithms: an overview," *npj Quantum Information*, vol. 2, Article 15023, 2016. (Nature)

8. M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2000. (Cambridge University Press & Assessment)

9. J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018. (arXiv)

10. S. Aaronson and A. Ambainis, "The need for structure in quantum speedups," *Theory of Computing*, vol. 10, no. 6, pp. 133–166, 2014 (arXiv earlier). (theoryofcomputing.org)

11. C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 623–656, July–October 1948. (Harvard Mathematics Department)

12. T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., Wiley, 2006. (Wiley Online Library)

13. R. M. Fano, *Transmission of Information: A Statistical Theory of Communications*, MIT Press, 1961 (source and classical statement of Fano's inequality). (Scholarpedia)

14. R. Landauer, "Irreversibility and heat generation in the computing process," *IBM Journal of Research and Development*, vol. 5, no. 3, pp. 183–191, 1961. (Worry Dream)

15. C. H. Bennett, "Logical reversibility of computation," *IBM Journal of Research and Development*, vol. 17, no. 6, pp. 525–532, 1973. (ACM Digital Library)

16. T. Toffoli, "Reversible computing," in *Automata, Languages and Programming* (Proc. ICALP), 1980 — classic account of reversible gates (Toffoli gate). (cqi.inf.usi.ch)

17. D. Boneh and V. Shoup, *A Graduate Course in Applied Cryptography*, Version (draft), Stanford University (on-line graduate textbook/draft widely used), 2020 (and later versions). (crypto.stanford.edu)

18. J. Katz and Y. Lindell, *Introduction to Modern Cryptography*, 2nd ed., CRC Press, 2014. (eclass.uniwa.gr)

19. S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof-systems," in *Proceedings of STOC*, 1985 — foundational zero-knowledge/knowledge complexity work. (MIT CSAIL People)

20. O. Goldreich, S. Micali, and A. Wigderson, "Proofs that yield nothing but their validity," *J. ACM*, vol. 38, no. 3, pp. 690–728, 1991 (GMW framework and constructions). (Institute for Advanced Study)

21. C. Gidney and M. Ekerå, "How to factor 2048-bit RSA integers in 8 hours using 20 million noisy qubits," *Quantum*, vol. 5, p. 433, 2021; arXiv:1905.09749 (detailed, modern resource estimates for Shor-type attacks). (Quantum)

22. C. Gidney and M. Ekerå, "How to factor 2048 bit RSA integers," *arXiv:1905.09749*, 2019 (preprint of the same resource-estimate work). (arXiv)

23. M. Grassl, B. Langenberg, M. Roetteler, and R. Steinwandt, "Applying Grover's algorithm to AES: quantum resource estimates," *arXiv:1512.04965*, 2015 (and associated conference versions). (arXiv)

24. M. Grassl et al., "Applying Grover's algorithm to AES: quantum resource estimates," in *Post-Quantum Cryptography* (LNCS), Springer, 2016 (conference version / book chapter). (ACM Digital Library)

25. M. Roetteler, M. Naehrig, K. M. Svore, and K. Lauter, "Quantum resource estimates for computing elliptic curve discrete logarithms," in *ASIACRYPT*, 2017 (detailed gate/qubit costs for Shor-style attacks on ECC). (arXiv)

26. J. Proos and C. Zalka, "Shor's discrete logarithm quantum algorithm for elliptic curves," *arXiv:quant-ph/0301141*, 2003 (implementation details and costs for ECC). (arXiv)

27. R. Van Meter and K. M. Itoh, "Fast quantum modular exponentiation," *Phys. Rev. A*, vol. 71, p. 052320, 2005; architecture-aware resource analysis for Shor's inner arithmetic. (Physical Review)

28. A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: towards practical large-scale quantum computation," *Phys. Rev. A*, vol. 86, p. 032324, 2012 (surface-code error correction and overhead estimates). (Physical Review)

29. D. Litinski, "A game of surface codes: large-scale fault-tolerant quantum computation with lattice surgery," *Quantum*, 2019 (space–time tradeoffs and practical compilation models for surface codes). (Quantum)

30. National Institute of Standards and Technology (NIST), "Post-Quantum Cryptography Standardization," NIST CSRC project page (overview of PQC and standardization process). (NIST Computer Security Resource Center)

31. NIST, *Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process*, NIST IR 8545 (and associated reports IR-series), 2024 (status & transition guidance). (NIST Technical Series)

32. S. Aaronson, "The limits of quantum computers," *Scientific American*, Mar. 2008 — non-technical perspective on when quantum speedups are and aren't likely. (cs.virginia.edu)

33. D. Gottesman, "Stabilizer codes and quantum error correction," Ph.D. thesis / arXiv:quant-ph/9705052, 1997 (stabilizer formalism and fault-tolerance foundations). (arXiv)

34. A. Ambainis, "Quantum lower bounds by quantum arguments," *arXiv:quant-ph/0002066*, 2000 — important lower-bound technique for quantum query complexity. (arXiv)

35. M. Boyer, G. Brassard, P. Høyer, and A. Tapp, "Tight bounds on quantum searching," *Fortschr. Phys.* / conference preprint (PhysComp drafts / arXiv), 1996–1998 (tight analysis of Grover and counting variants). (arXiv)

36. U.S. White House / OSTP, "Report on Post-Quantum Cryptography and Transition Planning" (overview / policy / government guidance on transition to quantum-safe crypto; 2024/2025 reports and briefs), and ancillary federal guidance documents. (The White House)

Appendix: Maximum Size of Computation Quantum/Classical

Notation & assumptions

- Alphabet size: $|\Sigma| = 72$ (30 lowercase + 30 uppercase + 12 digits).
- Secrets: six distinct secrets, each of length L. We examine $L = 6$ and $L = 12$.
- $U = 72^L$ — number of possible L-character strings.
- Private bijection count: $6! = 720$.
- We consider three combinatorial models (per L):
  1. **Distinct** (no repetition): $M_{\text{distinct}} = U(U-1)\cdots(U-5)$.

2. **Tight allow-repetitions** (unordered multiset + map): $M_{\text{rep,exact}} = 6!\binom{U+5}{6} = U(U+1)\cdots(U+5)$.
3. **Very loose** (ordered 6-tuple, repetition allowed, × map): $M_{\text{rep,loose}} = 6!\,U^6$.

- Per-round information bound: an observed leaf label yields at most $H_6 = \log_2 6 \approx 2.585$ bits.
- Rounds-to-uniqueness: $R_{\min} \gtrsim \log_6 |\mathcal{H}| = \dfrac{\log_2 |\mathcal{H}|}{\log_2 6}$.
- **Grover/BBBV**: quantum query lower bound for unstructured search is $\Theta(\sqrt{M})$ oracle calls for search space M. The usual "optimal" Grover iteration count is approximately $\left\lceil \frac{\pi}{4}\sqrt{M} \right\rceil$.
- Physical extreme assumptions used for wall-clock lower bounds:
  - Margolus–Levitin (ML) operation-rate per kg: $\nu_{\text{kg}} \approx 5.4256 \times 10^{50}$ ops/s/kg (used as an optimistic per-mass bound).
  - Baryonic mass of observable universe (used as an extreme available mass): $M_{\text{univ}} \approx 10^{53}$ kg (order of magnitude).
  - From these: $\nu_{\text{univ}} = \nu_{\text{kg}} \cdot M_{\text{univ}} \approx 5.4256 \times 10^{103}$ ops/s (extreme idealized instantaneous op rate).
  - Lloyd total-ops-over-history upper bound (used as a hard global cap): $\approx 10^{120}$ total elementary operations available in the universe over its history. (All these physical limits are standard estimates in the literature; we annotate sources below.)

Step-by-step derivation and arithmetic

1. Compute universe sizes U for $L = 6$ and $L = 12$.

1.1 Formula: $U = 72^L$.
1.2 For $L = 6$:
$U_6 = 72^6 = 139\,314\,069\,504$.
(Exactly $139, 314, 069, 504$.)
1.3 For $L = 12$:
$U_{12} = 72^{12} = (72^6)^2 = 139\,314\,069\,504^2$.
Computed exactly:
$U_{12} = 19\,408\,409\,961\,765\,342\,806\,016$.

(These exact integers were computed with arbitrary-precision arithmetic.)

---

2. Exact combinatorial counts (distinct, tight multiset, loose) for each L.

2A. $L = 6$ **exact integers**

2A.1 Distinct (no repetitions):

$$M_{\text{distinct},6} = U_6(U_6 - 1)\cdots(U_6 - 5).$$

Computed exact integer:

731088363577565404310584261068288872329465955062599633308300083200 0

(log10 $\approx$ 66.8639698714789)

2B. $L = 12$ **exact integers (big; shown as computed)**

2B.1 Distinct:

$$M_{\text{distinct},12} = U_{12}(U_{12} - 1)\cdots(U_{12} - 5).$$

Exact integer (134 digits — computed):

534490195473619995340666087941301726360703049132275668404460760756010886800636005516688865 58

(log10 $\approx$ 133.72793974305134)

> Note: For $U \gg 6$, rising/falling products differ negligibly relative to magnitude; hence the approximation $M \approx U^6$ is accurate in order-of-magnitude terms.

---

3. Approximate compact forms and log2 counts (useful for rounds)

3.1 Approximation: $|\mathcal{H}| \approx U^6 = 72^{6L} = 2^{6L\log_2 72}$.
Compute base-2 logarithms:

- For $L = 6$:

$$\log_2 |\mathcal{H}_6| \approx 6 \cdot 6 \cdot \log_2 72 = 36 \log_2 72 \approx 222.1 \text{ bits.}$$

  (Our code computed $\log_2 \approx 222.096\ldots$; we use 222.1.)

- For $L = 12$:

$$\log_2 |\mathcal{H}_{12}| \approx 72 \log_2 72 \approx 444.2 \text{ bits.}$$

3.2 Thus doubling L doubles the bit count, and squares $|\mathcal{H}|$: $|\mathcal{H}_{12}| \approx (|\mathcal{H}_6|)^2$.

---

4. Rounds-to-uniqueness $R_{\min}$

4.1 Per-round max leakage (half-open bound) $H_6 = \log_2 6 \approx 2.5849625$ bits.

4.2 Formula:

$$R_{\min} \gtrsim \frac{\log_2 |\mathcal{H}|}{\log_2 6}.$$

4.3 Compute:

- $L = 6$: $R_{\min,6} \approx \lceil 222.096/2.5849625 \rceil = \lceil 85.9 \rceil = 86$ rounds.
- $L = 12$: $R_{\min,12} \approx \lceil 444.192/2.5849625 \rceil = \lceil 171.9 \rceil = 172$ rounds.

Annotation: R doubles as L doubles.

---

5. Classical total-work lower bound (black-box)

5.1 Model: each candidate hypothesis must be checked (membership/consistency across R rounds). Cost per test ˜ linear in R. Lower bound:

$$W_{\text{class}}(L) = R_{\min}(L) \cdot |\mathcal{H}_L| \quad \text{(elementary ops, approximate)}.$$

5.2 Numeric approximations using the U^6 approximation:

- For $L = 6$:

  $$W_{\text{class},6} \approx 86 \times 7.3108836 \times 10^{66} \approx 6.29 \times 10^{68} \text{ elementary ops.}$$

  Our code produced $W_{\text{class},6} \approx 6.29 \times 10^{68}$ and log10 $\approx 68.799...$

- For $L = 12$:

  $$W_{\text{class},12} \approx 172 \times 5.34490195 \times 10^{133} \approx 9.19 \times 10^{135} \text{ elementary ops.}$$

  Code: $ \approx 9.19\times 10ˆ{135}$ (log10 $\approx 135.96...$).

Annotation: Doubling L multiplies the total-work by roughly $\frac{172 \cdot 72^{72}}{86 \cdot 72^{36}} \approx 2 \cdot 72^{36} \approx 2.0 \times |\mathcal{H}_6|$, i.e. about another $ \sim2ˆ{222}$ factor.

---

6. Quantum (Grover) lower bound: oracle calls and work

6.1 Best possible quantum scaling (unstructured): $\Theta(\sqrt{M})$ oracle calls for search space M.

6.2 The usual Grover iteration count (approx optimal):

$$\text{calls} \approx \left\lceil \frac{\pi}{4}\sqrt{M} \right\rceil.$$

6.3 Using $M \approx M_{\text{rep,exact}} \approx U^6$ (tight model):

- For $L = 6$: $\sqrt{M} \approx 2.7038646 \times 10^{33}$.
  Grover iterations $\approx (\pi/4)\sqrt{M} \approx 2.12361027 \times 10^{33}$ oracle calls.
- For $L = 12$: $\sqrt{M} \approx 7.310883636562886 \times 10^{66}$.
  Grover iterations $\approx (\pi/4)\sqrt{M} \approx 5.741954580968896 \times 10^{66}$ oracle calls.

Annotation: Squaring M (6→12) squares $\sqrt{M}$ and multiplies the quantum iteration count by $\approx 2^{111}$ (i.e., huge).

6.4 Note: each oracle call is not a single elementary op — it must compute the predicted labels across R rounds and compare to observed labels, revert ancillas, etc. So realistic per-oracle gate counts are large and multiply the total quantum gate count.

---

7. Physical upper bounds used (Margolus–Levitin, Lloyd, Bekenstein): numbers & sources

7.1 **Margolus–Levitin (ML)**: gives maximum rate of orthogonal transitions per energy: $\nu_{\max} = \dfrac{2E}{\pi\hbar}$. In operational terms, one can use derived numbers for per-kg limits. We use the conservative number:

$$\nu_{\text{kg}} \approx 5.4256 \times 10^{50} \text{ ops/s/kg}$$

(standard estimate used in "ultimate computer" style arguments — see Margolus & Levitin; Lloyd).

7.2 **Available mass**: baryonic mass of observable universe (order-of-magnitude):

$$M_{\text{univ}} \sim 10^{53} \text{ kg}.$$

7.3 Derived extreme instantaneous op-rate:

$$\nu_{\text{univ}} = \nu_{\text{kg}} \cdot M_{\text{univ}} \approx 5.4256 \times 10^{103} \text{ ops/s}.$$

(This is an extreme, physically permissible floor for instantaneous operations if you could convert all that mass into ML-limited processors and run them concurrently.)

7.4 **Lloyd total-ops bound** (integrated over cosmic history): approximate total elementary ops the universe can have performed $\sim 10^{120}$. (See S. Lloyd, "Ultimate physical limits to computation".)

7.5 **Bekenstein/Holographic bounds**: impose upper bounds on information density for a region of given energy/size (used conceptually to note that packing arbitrarily many bits into finite volume produces horizons/black holes).

(References: Margolus & Levitin; S. Lloyd; Bekenstein; these are the standard physics sources. I used these canonical numbers as the basis for the numeric floors above.)

---

8. Minimal wall-clock times under extreme ML/universe assumptions

*(These are **absolute lower bounds** obtained by dividing required elementary operations or oracle calls by ν_{univ}. They assume absurd engineering: full conversion of that mass into ML-limited processors, perfect reversibility/coherence, no overheads, no error correction cost, and that each required computational primitive maps to one ML-limited transition — i.e., the most optimistic physical floor. Use them only as theoretical lower bounds.)*

8.1 Compute minimal classical wall-clock (divide total classical elementary ops by ν_{univ}):

- For $L = 6$:

$$t_{\text{class,min},6} = \frac{W_{\text{class},6}}{\nu_{\text{univ}}} \approx \frac{6.29 \times 10^{68}}{5.4256 \times 10^{103}} \approx 1.159 \times 10^{-35} \text{ s.}$$

  (This is unrealistically tiny — because the assumptions convert everything into maximum-rate operations for a single unit time.)

- For $L = 12$:

$$t_{\text{class,min},12} = \frac{9.19 \times 10^{135}}{5.4256 \times 10^{103}} \approx 1.694 \times 10^{32} \text{ s} \approx 5.36 \times 10^{24} \text{ years.}$$

  (Huge; vastly exceeds the age of the universe.)

8.2 Compute minimal quantum wall-clock (divide Grover oracle-call count by ν_{univ}, again optimistic that one ML-op = one oracle call):

- For $L = 6$:

$$t_{\text{quant,min},6} \approx \frac{2.1236 \times 10^{33}}{5.4256 \times 10^{103}} \approx 3.914 \times 10^{-71} \text{ s.}$$

- For $L = 12$:

$$t_{\text{quant,min},12} \approx \frac{5.74195 \times 10^{66}}{5.4256 \times 10^{103}} \approx 1.0583 \times 10^{-37} \text{ s.}$$

Annotation: these times are formal lower bounds; they do **not** reflect the true cost of implementing Grover oracles (which require many physical gates, fault-tolerant overhead, and long coherent times). They demonstrate that the ML-derived instantaneous rate, when applied naively, gives tiny formal lower bounds — but those are *not* attainable in practice for such complex oracles.

---

9. Compare classical total-work to Lloyd universe total ops (feasibility over cosmic history)

9.1 Lloyd bound: total ops ever possible in our universe $\sim 10^{120}$.

9.2 Compare:

- $W_{\text{class},6} \approx 6.29 \times 10^{68} \ll 10^{120}$. So, **in principle** (counting raw total elementary-op budget over history), the universe could supply enough elementary ops aggregated across space/time to carry out a classical exhaustive search for $L = 6$. This is only a *counting* statement, not an engineering plan — it ignores distribution, coordination, memory, energy dissipation, gravity, horizon formation, etc.
- $W_{\text{class},12} \approx 9.19 \times 10^{135} \gg 10^{120}$. Thus the classical exhaustive search for $L = 12$ **exceeds** the entire universe's total-ops budget and is therefore **impossible even in principle** under Lloyd's accounting.

Conclusion: the regime change from $L = 6$ to $L = 12$ crosses the Lloyd feasibility threshold for classical brute force.

---

10. Incorporate realistic-oracle and fault-tolerance overheads (qualitative multiplier)

10.1 Real quantum or classical oracles are not single-primitive operations. They must:

- decode candidate secrets,
- for each of R rounds compute leaf membership (with modular rotations, table lookups or arithmetic),
- compare predicted label to observed label,
- restore ancillas (reversibility),
- manage error-correction for logical qubits (quantum) or manage I/O/VM/memory (classical).

10.2 Conservative estimate: let C_{oracle} denote elementary gates per oracle call. Plausible small values for minimal per-round reversible operations (per round) are tens to thousands; with R rounds, per-oracle cost $\sim c \cdot R$. For example:

- Take $c = 100$ primitive gates per round $\rightarrow$ per-oracle cost $\approx 100R$.
- For $L = 6$, $R \approx 86 \rightarrow$ per-oracle ≈ 8600 elementary ops.
- For $L = 12$, $R \approx 172 \rightarrow$ per-oracle $\approx 17{,}200$ elementary ops.

10.3 Therefore the **realistic** quantum total elementary-gate count ≈ grover_calls × per_oracle_cost, which multiplies the naive grover_calls-based wall-clock times above by many orders of magnitude (e.g., $10^{4-10}6$ or more), and the required number of physical qubits (for fault-tolerance) multiplies further (often by $10^{3-10}6$ in surface code estimates). So the formal tiny ML-derived times disappear under realistic cost modeling.

---

11. Final numeric summary (table of salient numbers — values from the calculations)

- $U_6 = 72^6 = 139\,314\,069\,504$.
- $U_{12} = 72^{12} = 19\,408\,409\,961\,765\,342\,806\,016$.
- Tight search space (approx):
  - $|\mathcal{H}_6| \approx 7.3108836 \times 10^{66}$.
  - $|\mathcal{H}_{12}| \approx 5.34490195 \times 10^{133}$.
- Rounds:
  - $R_{\min,6} \approx 86$.
  - $R_{\min,12} \approx 172$.
- Classical total-work (approx):
  - $W_{\text{class},6} \approx 6.29 \times 10^{68}$ elementary ops.
  - $W_{\text{class},12} \approx 9.19 \times 10^{135}$ elementary ops.
- Quantum Grover iterations (tight model; $(\pi/4)\sqrt{M}$):
  - $L = 6 : \approx 2.12 \times 10^{33}$ oracle calls.
  - $L = 12 : \approx 5.74 \times 10^{66}$ oracle calls.
- Universe-scale operational limits used:
  - $\nu_{\text{univ}} \approx 5.4256 \times 10^{103}$ ops/s (ML per-kg $\times 10^{53}$ kg).
  - Lloyd total-ops bound $\approx 10^{120}$ total ops over history.
- Minimal wall-clock (extreme ML-based lower bounds):
  - Classical min: $t_{\text{class,min},6} \approx 1.16 \times 10^{-35}$ s ; $t_{\text{class,min},12} \approx 1.69 \times 10^{32}$ s (˜ 5.36×10^{24} years).
  - Quantum min (oracle calls treated as single ML ops): $t_{\text{quant,min},6} \approx 3.91 \times 10^{-71}$ s ; $t_{\text{quant,min},12} \approx 1.06 \times 10^{-37}$ s.

- Feasibility under Lloyd:
  - Classical $L = 6$: **possible** in principle (W_class $\ll$ 10^120) — counting only.
  - Classical $L = 12$: **impossible** (W_class $\gg$ 10^120) — exceeds universe total-ops budget.

Important interpretation notes (closing)

- The **small ML-derived minimal times** for $L = 6$ are formal lower bounds that arise from dividing required ops by an extreme instantaneous operation rate; they are **not** achievable in practice because they ignore per-oracle gate counts, memory/IO, energy dissipation, decoherence, and general relativity/gravity constraints. They serve only to show absolute physical floors.
- The **Lloyd total-ops** bound ($\approx 10^{120}$) is a useful hard cap: if a required classical operation count exceeds it (as for $L = 12$), the attack is impossible even in principle over the entire life of the universe.
- The **quantum** attacker benefits from Grover's square-root speedup, but Grover cannot reduce an unstructured search to polynomial time; $\sqrt{|\mathcal{H}|}$ is still astronomically large for both $L = 6$ and $L = 12$, and realistic quantum resource multipliers (error correction, ancillas, T-count) make the needed resources effectively unreachable.

Sources and physical limits used

- Margolus, N. and Levitin, L. B., "The maximum speed of dynamical evolution," Physica D (1998) — Margolus–Levitin bound.
- S. Lloyd, "Ultimate physical limits to computation," Nature 406, 1047–1054 (2000) — Lloyd's total-ops and ultimate laptop/server calculations.
- R. Bousso, J. D. Bekenstein, and standard literature on Bekenstein/Holographic bounds (for the information-density constraints).
- Grover, Lov K., "A fast quantum mechanical algorithm for database search," STOC 1996; BBBV (Bennett et al.) and Zalka for quantum lower bounds/optimality.
- Quantum resource-estimate literature: Gidney & Ekerå, Grassl et al., Roetteler et al., and surface-code fault-tolerance surveys (for realistic multipliers on quantum gate counts).