DAC: Dual Asset Chain

A Quantum-Resistant Layer-1 Blockchain

Davide Costa
DAC Labs
davide.costa@dachain.tech

Fabio Fiori
DAC Labs
fabio.fiori@dachain.tech

Antonio La Gatta, Ugo Moschella
Q-Alliance, Lugano

Version 1.3
May 2026

Abstract

The imminent arrival of scalable quantum computers constitutes an existential threat to the cryptographic primitives that underpin all deployed public blockchain networks. RSA and elliptic-curve cryptography (ECC), which protect transaction signing and key exchange in Bitcoin and Ethereum, are fully broken by Shor's algorithm running on a fault-tolerant quantum computer with a few thousand logical qubits. SHA-256, the hash function driving classical Proof-of-Work mining, has its effective preimage resistance halved from 2^{256} to 2^{128} operations by Grover's algorithm, conferring a permanent quadratic mining advantage on any quantum miner.

This paper presents *DAC (Dual Asset Chain)*, a novel EVM-compatible Layer-1 blockchain that addresses the quantum threat at three independent levels. First, *Pure Entropy* $\mathcal{E}^*$ sources physically random bits from the quantum-mechanical Barkhausen effect in ferromagnetic materials, producing a true one-time-pad (OTP) with ≥ 0.99 bits of entropy per bit. Second, *Entropy-Padded Transactions (EPT)* XOR every transaction with a pre-assigned block of $\mathcal{E}^*$ before broadcast, guaranteeing information-theoretic confidentiality ($\mathcal{I}(\mathsf{TX}; \mathsf{CT}) = 0$) against any adversary—including an unbounded quantum attacker. Crucially, ECDSA transaction signing is *preserved without modification*: the OTP layer renders Shor's algorithm inapplicable, because a quantum adversary intercepting the broadcast observes only uniformly random ciphertext and never has access to the ECDSA public key or signature. Third, *Proof of Quantum Work (PoQW)* replaces SHA-256-based mining with quantum hashing executed on quantum annealers, making block production intrinsically quantum-intensive while reducing energy consumption by more than 99% compared to Bitcoin Proof-of-Work.

We formalise the security of each component, prove the unconditional secrecy of the EPT layer, characterise the quantum attack surfaces addressed, and compare DAC against existing post-quantum blockchain proposals. We also provide a rigorous justification for why neither lattice-based DSA (CRYSTALS-Dilithium) nor hash-based DSA (SPHINCS+, XMSS) is suitable for DAC's signing layer: lattice schemes embed hash functions as a structural requirement and rely on heuristic parameters whose tight quantum security proof remained incomplete until 2023 [1]; hash-based

stateless schemes exhibit concrete quantum chosen-message attacks reducing security from 2^{128} to 2^{43} [2]; and stateful hash-based schemes are operationally incompatible with distributed blockchain signing [3, 4]. An open challenge—the quantum-safe distribution of $\mathcal{E}^*$ to validator nodes—is stated precisely and identified as the central direction for future work.

Keywords: quantum-resistant blockchain, one-time pad, Barkhausen entropy, Proof of Quantum Work, Shor's algorithm, Grover's algorithm, EVM, post-quantum cryptography.

1 Introduction

The global financial, logistics, and identity infrastructure increasingly relies on public blockchain networks whose security ultimately rests on two computational hardness assumptions: the intractability of integer factorisation (underpinning RSA) and the intractability of the discrete logarithm problem over elliptic curves (underpinning ECDSA, the signature scheme used in both Bitcoin and Ethereum). Peter Shor's 1994 quantum algorithm [5] reduces both problems to polynomial time on a fault-tolerant quantum computer, rendering every existing blockchain signature scheme unconditionally insecure once sufficiently powerful quantum hardware becomes available.

The threat is no longer theoretical. Recent research [6,7] demonstrates that even *today's* quantum annealers can attack small elliptic curves via QUBO (Quadratic Unconstrained Binary Optimisation) formulations of the ECDLP, signalling that the transition period has begun.

Existing responses to this threat focus on replacing classical signature algorithms with post-quantum alternatives that introduce new computational hardness assumptions (lattice-based, hash-based, or code-based). Such substitutions trade one set of mathematical assumptions for another. They also leave untouched: (i) the confidentiality of transaction content during broadcast and mempool residence; and (ii) the quantum advantage in Proof-of-Work mining conferred by Grover's algorithm. DAC takes a different approach: rather than replacing ECDSA with a new signing scheme, it *shields* the existing signature from observation using an information-theoretic OTP layer, eliminating the attack surface entirely without introducing any new cryptographic assumptions.

DAC addresses all three dimensions simultaneously:

1. **Information-theoretic transaction confidentiality.** Each transaction is XORed with a block of pure quantum-sourced entropy $\mathcal{E}^*$ before broadcast. By Shannon's perfect secrecy theorem, the ciphertext is computationally and information-theoretically indistinguishable from random noise, regardless of the adversary's computational power.

2. **Quantum-native consensus.** Block mining is performed using a Proof of Quantum Work (PoQW) protocol [8] built on quantum annealing hardware, ensuring that only nodes with genuine quantum hardware can participate in mining.

3. **ECDSA preservation via OTP shielding.** Standard Ethereum ECDSA signing is retained unchanged. Because $\mathcal{I}(\mathsf{TX}; \mathsf{CT}) = 0$, no ECDSA public key or signature is ever visible during broadcast; Shor's algorithm has no observable target. No new cryptographic assumptions are introduced for the signing layer.

Contributions.

- A formal proof that EPT transactions achieve Shannon perfect secrecy independently of any computational hardness assumption (Theorem 7.3).
- A precise analysis of how Shor's and Grover's algorithms threaten RSA, ECC, and SHA-256 in blockchain contexts, citing concrete qubit estimates from the literature.
- The full DAC architecture integrating E*, EPT, and PoQW with the EVM execution environment.
- A formalisation of the *Entropy Pre-Assignment* model, which defines the invariants a valid entropy distribution scheme must satisfy. The concrete design of a quantum-safe distribution mechanism is identified as the central open problem and left for future work.

The remainder of this paper is organised as follows. Section 2 analyses the quantum threat to RSA, ECC, and SHA-256. Section 3 reviews the theoretical foundations. Section 5 describes the DAC architecture. Section 6 details the $\mathcal{E}^*$ hardware system. Section 7 defines Entropy-Padded Transactions. Section 8 covers Proof of Quantum Work. Section 9 traces the complete transaction lifecycle. Section 10 provides the security analysis. Section 12 concludes.

2 Quantum Threats to Classical Cryptography

2.1 RSA and Integer Factorisation

RSA [9] derives its security from the presumed intractability of factoring a large semi-prime $N = p \cdot q$ into its prime factors p and q. The fastest known classical algorithm, the General Number Field Sieve (GNFS), has sub-exponential but super-polynomial complexity:

$$T_{\text{GNFS}}(n) = O\big(\exp\big(c \cdot n^{1/3}(\log n)^{2/3}\big)\big)\,, \tag{1}$$

where $n = \log_2 N$ is the bit-length of the modulus and $c \approx 1.923$.

Shor's quantum algorithm [5] solves the integer factorisation problem in *polynomial* time. On a gate-model quantum computer with $O(n)$ logical qubits and $O(n^3)$ gate operations, Shor reduces the factoring problem to order-finding via the quantum Fourier transform. Beauregard [10] showed that factoring an n-bit modulus requires approximately $2n$ logical qubits.

For a 2048-bit RSA key (the current minimum recommendation):

- Qubits required: ≈ 4096 logical qubits.
- Gate count: $O((2048)^3) \approx 8.6 \times 10^9$ Toffoli gates.
- Estimated time on a fault-tolerant quantum computer: ∼hours to days depending on physical error rate and code distance.

Remark 2.1. *While blockchains do not use RSA for on-chain signatures, RSA and similar integer-factorisation schemes secure the underlying P2P network routing, node communication, and legacy TLS layers. Furthermore, as shown below, the elliptic-curve cryptography that* does *secure the blockchain is actually more vulnerable to Shor's algorithm than RSA.*

2.2 Elliptic Curve Cryptography

Elliptic curve cryptography bases its security on the *Elliptic Curve Discrete Logarithm Problem* (ECDLP): given points P and $Q = kP$ on an elliptic curve E over a finite field $\mathbb{F}_p$, find the scalar $k \in \mathbb{Z}$. The best known classical algorithm for generic curves (Pollard's ρ) has fully exponential complexity $O(\sqrt{|\langle P \rangle|})$, which is why 256-bit ECC keys (as used in Bitcoin's secp256k1 and Ethereum's accounts) achieve security comparable to 3072-bit RSA.

Shor's discrete logarithm algorithm [5], adapted for elliptic curves by Proos and Zalka [6], solves the ECDLP in $O(n^3)$ time using $O(n)$ logical qubits. Their detailed resource analysis establishes:

Theorem 2.2 (Proos–Zalka qubit estimate [6])**.** *Shor's ECDLP algorithm over $\mathbb{F}_p$ with an n-bit prime p requires approximately $6n$ logical qubits and a gate depth of $O(n^2)$ for the core arithmetic (modular inversion via the extended Euclidean algorithm).*

For Ethereum/Bitcoin's secp256k1 curve ($n = 256$):

$$\text{Qubits} \approx 6 \times 256 = 1536 \text{ logical qubits.} \tag{2}$$

This establishes a critical vulnerability: the ECDLP requires substantially fewer qubits to break than equivalent RSA security.

Quantum annealing attack. Beyond gate-model quantum computers, recent work explores ECDLP attacks using *quantum annealing*. Dzierzkowski [7] shows how to convert the ECDLP into a QUBO problem solvable on quantum annealing hardware. The key reduction maps the ECDLP over a short Weierstraß curve to a system of polynomial equations over $\mathbb{F}_2$, which is then expressed as an Ising Hamiltonian:

$$H_{\text{Ising}} = -\sum_{i<j} J_{ij}\, \sigma_i \sigma_j - \sum_i h_i\, \sigma_i, \quad \sigma_i \in \{-1, +1\}. \tag{3}$$

Experimental results on curves of order 7 over 4-bit fields confirm the mathematical feasibility. However, scaling this to standard 256-bit sizes presents a steep structural hurdle. Transforming massive prime-field operations like modular inversion into a QUBO function causes a major explosion of auxiliary variables and heavily interconnected graphs. Because modern physical annealers rely on localized, sparse hardware topology, mapping a full secp256k1 curve remains a profound physical challenge. Therefore, while gate-model quantum computers running Shor's algorithm remain the primary operational threat vector we must defend against these optimization methods are theoretically fascinating and represents a long-term alternative threat to gate-model architectures.

Implication for Ethereum and Bitcoin. Every Ethereum account and UTXO in Bitcoin is secured by a secp256k1 key pair. Once the public key is exposed (which occurs the first time a user signs a transaction), a quantum adversary can recover the private key via Shor's ECDLP algorithm and forge arbitrary transactions. This places all previously-published addresses at risk.

2.3 Cryptographic Hash Functions and SHA-256

Hash functions do not admit polynomial-time quantum attacks analogous to Shor's algorithm. However, Grover's algorithm [11] provides a *quadratic* speedup for unstructured search, with significant consequences for blockchain mechanics.

Preimage resistance. For a hash function $H : \{0,1\}^* \to \{0,1\}^n$ with n-bit output:

$$T_{\text{classical}} = O(2^n), \tag{4}$$

$$T_{\text{quantum}} = O(2^{n/2}) \quad \text{(Grover's algorithm)}. \tag{5}$$

SHA-256 ($n = 256$): quantum preimage resistance drops from 2^{256} to 2^{128}.

Collision resistance. Classical birthday bound gives $O(2^{n/2})$ collisions; the quantum Brassard–Høyer–Tapp (BHT) algorithm [12] achieves:

$$T_{\text{BHT}} = O(2^{n/3}). \tag{6}$$

For SHA-256: quantum collision resistance is $2^{85.3}$ operations. While achieving this in practice requires an exceptionally large and highly coherent Quantum RAM (qRAM) architecture that remains beyond near-term hardware, it constitutes a permanent theoretical vulnerability.

Mining advantage and the hardware arms race. Bitcoin's Proof-of-Work requires miners to find a nonce x such that $H(B\|x) < \tau$ for a target threshold τ. A classical miner runs an expected $T = 2^{256}/\tau$ trials. A Grover miner evaluates the space with an algorithmic query complexity of $O(\sqrt{2^{256}/\tau})$.

While practical extraction of this advantage is currently bottlenecked by the strict limits of quantum circuit depth and the inability of Grover's algorithm to parallelize efficiently compared to classical ASICs, the fundamental tension between brute-force classical energy consumption and quantum query reduction remains. This architectural mismatch strongly motivates a transition away from SHA-256 PoW to protocols that anchor consensus in native quantum hardware (like DAC's PoQW), thereby neutralizing the algorithmic race while drastically lowering energy overhead.

Hash functions in DAC. DAC's core EPT layer and signing path use *no hash function*. The entropy identifier is a plain integer index j; transaction encryption is a one-time pad; signing is standard ECDSA.

Remark 2.3. *Transaction encryption uses a one-time pad ($\mathsf{CT} = \mathsf{TX} \oplus \mathcal{E}^*$), whose security is information-theoretic and entirely independent of any hash function or Grover's algorithm. See Theorem 7.3.*

2.4 Combined Impact on Blockchain Infrastructure

Table 1 summarises the quantum attack surface of existing blockchain cryptography.

Table 1: Quantum attack surface of existing blockchain cryptography

Primitive	**Blockchain use**	**Quantum algorithm**
RSA-2048	Legacy P2P / TLS	Shor
ECDSA-256	TX signatures	Shor
secp256k1	Wallet addresses	Shor+QA
SHA-256	PoW mining	Grover

Without mitigation, ECDSA exposure allows transaction forgery while classical miners are forced into an unsustainable parallelization race against Grover's algorithm. DAC addresses both attack surfaces through architecturally distinct mechanisms that introduce no new hash or lattice-based assumptions.

3 Theoretical Foundations

3.1 Shannon's Perfect Secrecy

Definition 3.1 (Shannon entropy [13])**.** *Let X be a discrete random variable over alphabet $\mathcal{X}$. The Shannon entropy of X is*

$$H(X) = -\sum_{x \in \mathcal{X}} \Pr[X = x] \log_2 \Pr[X = x]. \tag{7}$$

Definition 3.2 (Perfect secrecy [13])**.** *An encryption scheme* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *over message space $\mathcal{M}$ and ciphertext space $\mathcal{C}$ satisfies* perfect secrecy *if, for every distribution over $\mathcal{M}$, every $m \in \mathcal{M}$, and every $c \in \mathcal{C}$:*

$$\Pr[\mathsf{TX} = m \mid \mathsf{CT} = c] = \Pr[\mathsf{TX} = m]. \tag{8}$$

Equivalently, the mutual information between message and ciphertext is zero:

$$\mathcal{I}(\mathsf{TX}\,;\,\mathsf{CT}) = 0. \tag{9}$$

Theorem 3.3 (One-Time Pad perfect secrecy [13])**.** *Let $\mathcal{E}^* \in \{0,1\}^\ell$ be drawn uniformly at random and independent of the message, and used exactly once. Then $\mathsf{CT} = \mathsf{TX} \oplus \mathcal{E}^*$ satisfies perfect secrecy for all $\mathsf{TX} \in \{0,1\}^\ell$.*

Proof. Established by Shannon [13]; follows directly from Bayes' theorem showing $\Pr[\mathsf{CT} = c]$ is uniform and independent of TX. □

Corollary 3.4 (Quantum immunity of OTP)**.** *No quantum adversary, including one equipped with Grover's algorithm or Shor's algorithm, can distinguish* CT *from a uniformly random string or recover any partial information about* TX *from* CT *alone.*

Proof. Perfect secrecy is an information-theoretic property. It holds against adversaries with unbounded computational resources, including quantum computers. No algorithm—classical or quantum—can extract information that does not exist in the probability distribution. Since $\mathcal{I}(\mathsf{TX}; \mathsf{CT}) = 0$, no observable function of CT is correlated with TX. □

3.2 The Barkhausen Effect as a Quantum Entropy Source

A true OTP requires a physically random entropy source—one whose output cannot be predicted by any computational process. DAC uses the Barkhausen effect, a quantum-mechanical phenomenon in ferromagnetic materials, as its primary entropy source.

Physical mechanism. When a ferromagnetic material is subjected to a slowly varying external magnetic field $H(t)$, the magnetisation M does not change continuously but instead exhibits discrete, irreversible jumps as magnetic domain walls de-pin from lattice impurities and re-anchor at new metastable positions. These jumps are the Barkhausen noise signal.

Thermal vs. quantum regime. The microscopic origin of these jumps depends on temperature. Above the classical/quantum crossover, domain-wall depinning is dominated by thermal activation over local energy barriers, and the resulting noise is high-entropy but classical-stochastic in character. Deep in the quantum regime—reached in suitable Ising magnets at sub-Kelvin temperatures—domain-wall motion is instead initiated by quantum tunnelling of correlated spin plaquettes through the same barriers. Simon, Silevitch, Stamp, and Rosenbaum established this directly in $LiHo_{0.4}Y_{0.6}F_4$ at $T \in [90, 580]$ mK, observing two distinct quantum tunnelling mechanisms (independent and cooperative domain-wall cotunnelling) with avalanches of 1.5×10^{15} to 4.5×10^{16} spins per event and a temperature-independent event statistics that excludes thermal activation as the driver. This result is the experimental foundation for treating cryogenically operated Barkhausen sources as quantum-mechanical entropy sources in the strong sense, and motivates the hybrid-temperature architecture described below.

Macroscopic dynamics: Classical vs. Quantum regimes. In the classical thermal regime, the discontinuous domain-wall depinning is correctly governed by the classical Alessandro-Beatrice-Bertotti-Montorsi (ABBM) model, which treats the domain wall as an elastic interface driven by thermal activation through a random pinning potential.

However, deep in the sub-Kelvin regime, this semi-classical Langevin-type interface model becomes fundamentally insufficient. Deprived of the thermal energy required to overcome local pinning barriers, domain-wall depinning must instead be modelled as a macroscopic quantum tunnelling (MQT) event. In this quantum Barkhausen regime, the transition rate is determined by an effective instanton bounce action S_B through the pinning potential, driven entirely by the quantum cotunnelling of correlated spin plaquettes rather than classical thermal fluctuations.

Statistical signature. Jump amplitudes s follow a power-law distribution characteristic of a system at self-organised criticality:

$$P(s) \propto s^{-\tau}, \quad \tau \approx 1.5, \tag{10}$$

and the energy of individual events scales as $P(E) \propto E^{-\epsilon}$ with $\epsilon \approx 2$.

Physically, these Barkhausen jumps cluster in time, creating highly correlated avalanches. Therefore, the raw physical signal possesses strong short-term temporal correlations. Because von Neumann unbiasing cannot break these temporal correlations, the digitised raw signal must be passed through a robust cryptographic randomness extractor (such as a universal hash function or SP 800-90A conditioning component, detailed in Section 6). This guarantees the resulting cryptographic bitstream $\mathcal{E}^* \in \{0,1\}^*$ achieves strict statistical independence. Because the variables are binary $\{0,1\}$, the autocovariance of this post-processed stream satisfies:

$$C_{\mathcal{E}}(k) = \mathbb{E}[(\mathcal{E}^*_i - \mu)(\mathcal{E}^*_{i+k} - \mu)] \approx 0 \quad \text{for } k \neq 0, \tag{11}$$

where $\mu = 0.5$, confirming the complete absence of any exploitable temporal structure in the final one-time pad.

Hybrid two-tier entropy architecture. DAC operates two classes of $\mathcal{E}^*$ source distinguished by their thermal regime:

- **Server-side (cryogenic) sources.** Validator-class $\mathcal{E}^*$ modules are operated at sub-Kelvin temperatures using helium-dilution refrigeration, placing the Barkhausen source deep in the quantum regime. These sources are the canonical $\mathcal{E}^*$ stream from which OTP material is minted.

- **Local (room-temperature) sources.** Client-side devices operate at ambient temperature for cost and scalability reasons. In this regime Barkhausen jumps are thermally activated and the signal is statistically high-entropy rather than quantum-pure. Local entropy is never used in isolation: it is XOR-mixed with a fresh server-minted block before any OTP role, so its contribution is conditioned on the cryogenic stream.

Formally, to eliminate the architectural disconnect between generation and consumption, the unified one-time pad applied by the client during transaction encryption (Section 7) is constructed as:

$$\mathcal{E}^*_{\text{final}} = \mathcal{E}^*_{\text{server}} \oplus \mathcal{E}^*_{\text{local}} \tag{12}$$

This guarantees the encryption achieves security at least equivalent to the purely quantum cryogenic stream, while successfully integrating local device stochasticity.

At both tiers, the Barkhausen source is supplemented with two additional physical processes for redundancy and whitening:

1. **Johnson-Nyquist thermal noise.** A resistor R at temperature T generates voltage fluctuations with power spectral density $S_V(f) = 4k_BTR$, contributing $\sigma^2 = 4k_BTR\Delta f$ per bandwidth Δf.

2. **Avalanche noise.** A reverse-biased Zener diode operating in avalanche breakdown generates shot noise from stochastic impact ionisation events.

The three sources are combined analogically before digitisation, resulting in a broadband, high-variance physical noise signal with a near-Gaussian amplitude distribution. While this analog mixing guarantees a physically unpredictable entropy source, raw digitised Gaussian noise does not inherently achieve perfect cryptographic uniformity. To bridge this gap, the raw digitised bitstream is passed through a strict cryptographic conditioning phase (including a von Neumann decorrelator and an optional SP 800-90A Hash DRBG, detailed in Section 6.1). It is only after this digital extraction phase that the final combined entropy rate satisfies $H(\mathcal{E}^*_i) \geq (1-\epsilon) \cdot 1$ bit/bit with $\epsilon < 0.01$ under NIST SP 800-90B testing.

Remark 3.5 ("Quantum" as physical implementation, not security prerequisite)**.** *The OTP-based confidentiality result for EPT (Theorem 7.3) is purely information-theoretic: it follows from $\mathcal{I}(TX; CT) = 0$ whenever the pad is uniformly distributed and independent of TX, with no further assumption on the physical origin of the pad. The qualifier "quantum" applied to $\mathcal{E}^*$ in this paper refers to the physical implementation of the entropy source—in particular to the cryogenic Barkhausen mechanism—not to a theoretical prerequisite of the security argument. Equivalently: a thermally activated Barkhausen source that is independently certified to satisfy Lemma 6.1 would yield the same secrecy bound. The motivation for cryogenic operation is therefore to obtain a source whose entropy quality is anchored in fundamental physics rather than in environmental modelling, and whose independence properties are verifiable from first principles.*

3.3 Quantum Annealing

Quantum annealing [14] is a metaheuristic that exploits quantum tunnelling and superposition to minimise an Ising Hamiltonian:

$$H = -\sum_{i<j} J_{ij}\,\sigma_i\sigma_j - \sum_i h_i\,\sigma_i, \quad \sigma_i \in \{-1,+1\}. \tag{13}$$

Any NP-hard combinatorial problem can be formulated as the minimisation of a Quadratic Unconstrained Binary Optimisation (QUBO) objective:

$$\min_{\mathbf{x}\in\{0,1\}^n} \mathbf{x}^\top Q\,\mathbf{x}, \tag{14}$$

which maps directly to the Ising model via $x_i = (\sigma_i + 1)/2$. Quantum annealers implement this directly in hardware, using programmable superconducting qubit couplers. Current-generation processors provide > 5000 physical qubits and > 35000 couplers, enabling the solution of QUBO instances intractable on classical hardware for certain problem structures [8]. DAC's PoQW protocol (Section 8) exploits this capability for consensus.

4 On the Choice of Signing Architecture

A natural response to the quantum threat against ECDSA is to replace it with a *post-quantum* digital signature scheme. Two families dominate the NIST PQC standardisation landscape: lattice-based schemes (principally CRYSTALS-Dilithium / ML-DSA [15]) and hash-based schemes (SPHINCS+ [16], XMSS [17]).[1] This section evaluates both families against DAC's requirements and explains why neither is adopted for the transaction-signing layer. The conclusion is that the EPT layer, by guaranteeing $I(\mathsf{vk};\mathsf{CT}) = 0$, renders Shor's algorithm inapplicable and eliminates the need for any alternative signing infrastructure.

4.1 Lattice-Based DSA: Architectural Hash Dependency

DAC's core security requirement for the signing layer is that it must not introduce any dependency on the computational hardness of hash functions. CRYSTALS-Dilithium [19], the primary NIST-selected lattice-based signature scheme (standardised as FIPS 204 [15]), fails this requirement by design. Dilithium's signing algorithm generates a challenge polynomial as

$$c \;=\; H(M \,\|\, \mathbf{w}_1), \tag{15}$$

where $H : \{0,1\}^* \to B_\tau$ is instantiated as `SHAKE-256` for security levels 2 and 3, and the matrix $\mathbf{A}$ used in every key generation and verification is expanded from a seed via `SHAKE-128`. As noted in the original specification [19], *"the two operations that constitute nearly the entirety of the signing and verification procedures are expansion of an XOF (we use SHAKE-128 and SHAKE-256)"*. The challenge hash in Eq. (15) is not an optimisation but a *security-critical* component: removing it would break the zero-knowledge property of the underlying Fiat-Shamir transform and invalidate all known security proofs [1, 19].

[1] A third post-quantum family—isogeny-based signatures—lies outside our scope. The state of the art, SQIsignHD [18], achieves the most compact post-quantum signatures known (109 bytes at NIST-1), but its verification cost is estimated at $\gtrsim 200$ ms in optimised form (600 ms in the authors' proof-of-concept `sagemath` implementation), its security reduction relies on the same isogeny-in-higher-dimension machinery that broke SIDH in 2022–2023, and no isogeny scheme has entered the NIST PQC finalisation track. These properties disqualify it from DAC's validator signing layer regardless of the signature-size advantage.

Remark 4.1 (Dilithium requires hash functions unconditionally)**.** *There is no known variant of Dilithium that eliminates the hash-function dependency. The SHAKE calls are architecturally inseparable from the protocol: they appear in key generation (seed expansion of* $\mathbf{A}$*), signing (challenge generation, commitment hashing), and verification (recomputing the challenge). Adopting Dilithium would therefore directly contradict DAC's requirement that the signing layer be hash-function-free.*

A broader survey of lattice-based DSA confirms this pattern: every standardised or near-standardised lattice signature scheme—including Falcon [15] and the general Hash-and-Sign paradigm—embeds a cryptographic hash function as a structural primitive [20]. This is not a limitation of current implementations; it is a consequence of how the Fiat-Shamir transform with aborts achieves its security proof in the (quantum) random oracle model.

4.2 Lattice-Based DSA: Incomplete Quantum Security Proof

Even setting aside the hash-function dependency, the quantum security of Dilithium rests on foundations that were incompletely understood until recently. Dilithium's security is reduced to the hardness of three computational problems: Module Learning With Errors (MLWE), Module Short Integer Solution (MSIS), and *SelfTargetMSIS*. The first two are well-studied; the third is a non-standard variant introduced specifically for this proof. As Jackson, Miller, and Wang established in 2023 [1], *the quantum hardness of SelfTargetMSIS was an open problem* until their work. Prior to that result, the only rigorous quantum security proof for Dilithium [19] required modifying the modulus to $q \equiv 5 \pmod 8$—a condition *incompatible* with the original Dilithium specification, which requires $q \equiv 1 \pmod{2n}$ to enable the Number Theoretic Transform (NTT) that underpins its efficiency. Jackson et al. provided the first proof of SelfTargetMSIS hardness in the QROM via a reduction from MLWE, but at the cost of new parameter sets. Concretely, their provably secure parameters require:

- Public-key sizes $\approx 11.4\times$ larger than standard Dilithium [1].
- Signature sizes $\approx 3.2\times$ larger than standard Dilithium [1].

The NIST-standardised parameters in FIPS 204 are therefore *heuristically chosen*: they optimise for efficiency at the expense of tight quantum security reductions. While relying on heuristic parameter selection is a standard and pragmatic choice in applied cryptography, it embeds computational assumptions that DAC's information-theoretic architecture specifically seeks to avoid. A blockchain deploying Dilithium at provably secure parameters (thereby eliminating the heuristic gap) would face an 11-fold increase in public-key bandwidth—an unacceptable overhead for a high-throughput chain [1, 20].

4.3 Hash-Based DSA: Quantum Chosen-Message Attacks

Hash-based signature schemes are attractive because their security reduces to hash-function properties rather than algebraic hard problems. The leading stateless candidate, SPHINCS+ [16], was standardised by NIST alongside Dilithium. However, its security model does not cover the *quantum-access* threat relevant to blockchain validators. Yuan, Tibouchi, and Abe [2] demonstrated concrete quantum chosen-message attacks (qCMA) on both SPHINCS and SPHINCS+ under the *quantum-access* (Q2) security model, in which

the adversary can query the signing oracle in superposition. Their results are summarised in Table 2.

Table 2: Quantum chosen-message attacks on SPHINCS(+) [2]. q_s = quantum signing queries; q_H = quantum hash queries. Classical security is the advertised EUF-CMA bound; attack complexity is the demonstrated qCMA upper bound.

Scheme	Classical	qCMA Attack (PO model)	
	$\log_2 q_H$	$\log_2 q_s$	$\log_2 q_H$
SPHINCS-256	128	43	43
SPHINCS+-256s	128	48	80
SPHINCS+-256f	128	46	80

The core mechanism of the attack is a Grover search over the FORS index selection function. By making $O(2^{h/2})$ quantum signing queries, the adversary finds enough few-time signatures associated with a single index to recover the corresponding FORS secret key, enabling unlimited forgery.

It is crucial to acknowledge that current blockchain networks operate over classical TCP/IP protocols. Therefore, a remote adversary cannot physically send a quantum state (a superposition of queries) to a classical validator today. Under this strict classical-access (Q1) model, SPHINCS+ remains secure against the qCMA attack.

However, DAC is designed as a foundational, long-term Layer-1 architecture. As global communication infrastructure inevitably evolves toward quantum networking, or as validators are deployed in hybrid quantum-classical hardware enclaves, the physical barrier preventing Q2 access will dissolve. Relying on the classicality of the transport layer is a temporary, brittle defense. To guarantee permanent safety, a truly post-quantum blockchain must assume the strongest possible adversary—one capable of superposition queries (the Q2 model). Under this model, SPHINCS+ provides at most 2^{43}–2^{48} quantum signing queries of security—far below the claimed 2^{128} post-quantum level [2].

4.4 Hash-Based DSA: State Management in Distributed Consensus

Stateful hash-based signature schemes—XMSS, LMS, HSS—avoid the qCMA vulnerability by using each one-time signature (OTS) key at most once, tracked via a monotonically incrementing counter. This per-sign state update is the source of their security *and* their primary deployment hazard. McGrew et al. [3] and the 2024 IETF draft by Wiggers et al. [4] identify the following failure modes, each of which allows an adversary to forge signatures:

1. **State synchronisation failure.** If the signing process crashes between signing and persisting the updated counter to non-volatile storage, the same OTS key is reused on the next restart. As noted in [4]: *"double-signing with the same OTS key allows forgeries"* and "extreme care should be taken to avoid the risk that an OTS key will be reused accidentally."

2. **Non-volatile cloning.** Backup restoration, VM image cloning, or node migration all reset the counter to a previously-used state. On a blockchain, node operators routinely restore from snapshots.

3. **Concurrent signing.** Multiple validator instances sharing duties cannot atomically claim disjoint OTS key ranges without a coordination protocol. No universal solution exists for distributed environments [4].

The stateless alternative SPHINCS+ avoids the state problem but at the cost of large signatures: $\approx$ 41 KB for 128-bit security at 2^{60} messages (Table 1 of [3]), compared to 64 bytes for ECDSA. At 1000 transactions per second, SPHINCS+ signatures alone would require ≈ 3.3 GB/hour of additional network bandwidth per validator, an order-of-magnitude increase over ECDSA. Furthermore, the multi-target quantum attack [17] compounds the qCMA problem at blockchain scale. When the same hash function family is applied to d targets—as occurs in a Merkle tree over 2^h OTS keys—the quantum attack complexity drops from $O(2^n)$ to $O(2^n/\sqrt{d})$. For XMSS with tree height $h = 60$ ($d \approx 2^{66}$ hash images), achieving 256-bit quantum security requires a hash output of 322 bits rather than 256 [17], directly increasing signature sizes by a further 25%.

Remark 4.2 (Deployed instance: Quip.Network)**.** *A concrete deployed example of the constraints above is the Quip.Network protocol [21], which adopts* raw WOTS+ *over Keccak-256 as its on-chain signing primitive. Because no Merkle or hypertree structure is layered above the one-time key, every transaction consumes the signer's entire WOTS+ key and any change must be routed to a freshly generated QUIP address—a requirement of the underlying scheme rather than a design choice. This forced per-spend address rotation is incompatible with the EVM account model, in which an externally-owned account is a long-lived* $(\mathsf{sk}, \mathsf{vk})$ *pair bound to contract storage, allowances, and on-chain identity. DAC's rejection of a native hash-based signing layer in favour of the information-theoretic construction of Section 4.5 therefore aligns with the observation that even the simplest hash-based deployment pays its security tax in account-model and wallet UX obligations—not merely in signature size.*

4.5 DAC's Resolution: Information-Theoretic Signing Security

The analysis above shows that both major families of post-quantum DSA introduce unacceptable trade-offs for DAC's signing layer:

- **Lattice DSA** (Dilithium / FIPS 204) embeds hash functions as a structural requirement and relies on heuristically chosen parameters whose tight quantum security was open until 2023 [1].
- **Hash-based DSA** (SPHINCS+, XMSS) either suffers from demonstrated qCMA attacks reducing security from 2^{128} to 2^{43} [2], or introduces OTS state management hazards incompatible with distributed blockchain operations [3, 4].

DAC's EPT layer resolves the signing security problem by a different mechanism: it eliminates the *observation* of the signing key rather than attempting to harden the signing algorithm. Because $\mathsf{CT} = \mathsf{TX} \oplus \mathcal{E}_j^*$ and $I(\mathsf{TX}; \mathsf{CT}) = 0$ (Theorem 7.3), the ECDSA public key $\mathsf{vk}_{\mathsf{Alice}}$—which is embedded in TX—is information-theoretically hidden from any adversary who observes the broadcast tuple $(\mathsf{CT}, j, \mathsf{sig})$. Shor's algorithm requires the public key as input; it cannot operate on a quantity it cannot observe. This provides a strictly *stronger* guarantee than any computational post-quantum assumption:

- No hardness assumption is required—the protection is unconditional.

- No hash functions are introduced into the signing path.
- ECDSA is preserved unchanged, ensuring full EVM compatibility and zero migration cost for existing Ethereum tooling.
- Signature sizes remain 64 bytes (secp256k1 compact form).

The residual exposure window—after $\mathcal{E}_j^*$ is revealed post-finalization and TX (including $\mathsf{vk}_{\mathsf{Alice}}$) becomes public—is addressed by an address hygiene recommendation (Section 7): one fresh address per transaction ensures that the revealed public key governs an already-spent account. No long-term key exposure remains.

5 DAC Architecture

5.1 Design Goals

DAC is an EVM-compatible Layer-1 blockchain designed with the following security objectives:

1. **Quantum resistance of transaction confidentiality:** information-theoretic via OTP encryption; no computational assumption.
2. **ECDSA signing preserved:** standard Ethereum ECDSA retained; the OTP layer prevents any quantum attacker from observing the public key.
3. **Quantum-native mining:** PoQW makes block production classically intractable.
4. **EVM compatibility:** full support for Ethereum smart contracts and tooling.

5.2 Three-Layer Hierarchical Architecture

DAC organises network functionality into three layers:

SyncroChain The top-layer consensus chain. Supervisory Nodes (SNs) participate in epoch sealing, checkpoint creation, and randomness accumulation (RANDAO-style). The SyncroChain maintains global finality and entropy commitment records.

MasterChain The middle aggregation layer. Validator Nodes (VNs) collect shard-level block headers, validate cross-shard receipts, and coordinate entropy pre-assignment records.

ShardChains Multiple parallel execution shards running the EVM. Each shard has its own committee of VNs and processes a subset of pending encrypted transactions.

5.3 Node Types and Stake Requirements

Every Validator Node is equipped with an $\mathcal{E}^*$ hardware module (Section 6) and maintains a local encrypted entropy store indexed by integer entropy serial indices j.

Table 3: DAC node types and requirements

Type	Stake (DACT)	Role	Entropy store
Supervisory	100,000	Epoch seal, backup	Full pool copy
Validator	1,000	Block proposal, EPT	Pre-assigned $\mathcal{E}^*$
Light	0	Query, relay	None

5.4 Token Economy

DAC employs a dual-token model:

DACT Governance and staking token. Fixed supply of 10^9 tokens. Required to operate Supervisory and Validator Nodes.

DACC Transaction and reward coin. Dynamically minted as block rewards; burned on transaction fee payment, creating a deflationary equilibrium.

6 The Pure Entropy $\mathcal{E}^*$ System

6.1 Hardware Architecture

The $\mathcal{E}^*$ system comprises four physical components:

1. **Quantum Entropy Generator (QEG).** A ferromagnetic pickup coil wound around a soft-iron rod, driven by a slow triangle-wave excitation field. The Barkhausen pulses are amplified and band-pass filtered to extract the entropy-bearing component.
2. **Auxiliary noise sources.** Thermal resistor and avalanche diode, combined with the QEG output in an analogue mixer.
3. **Post-processing stage.** ADC at ≥ 1 MSps, followed by a von Neumann decorrelator and optional SP 800-90A Hash_DRBG for conditioning.
4. $\mathcal{E}^*$ **Dongle.** A tamper-resistant USB cryptographic module that stores the processed entropy in a key lifecycle manager, presenting blocks of $\mathcal{E}^*$ to applications via a standardised API.

The combined output satisfies:

- Entropy per bit: $H(\mathcal{E}_i^*) \geq 0.99$ bits/bit.
- NIST SP 800-90B min-entropy estimate: $\hat{H}_{\min} \geq 0.98$ bits/bit.
- Frequency test: bit balance $|f_1 - 0.5| < 0.02$.
- Autocorrelation: $|R_{\mathcal{E}}^*(k)| < 0.02$ for all $k \geq 1$.
- Serial compression test: incompressible at ≤ 1.01 bits/bit.

6.2 Formal Security Proof Against Quantum Attacks

Lemma 6.1 (Hardware entropy certification)**.** *The $\mathcal{E}^*$ hardware module is certified to satisfy the measurable randomness criteria of [22], §2.6:*

1. *per-bit Shannon entropy $H(\mathcal{E}^{*(i)}_j) \geq 0.99$ bits/bit (where $\mathcal{E}^{*(i)}_j$ is the ith bit of block $\mathcal{E}^*_j$), estimated via the NIST SP 800-90B suite [22, 23];*
2. *autocorrelation $|R^*_{\mathcal{E}}(k)| < 0.02$ for all lags $k \geq 1$;*
3. *frequency balance $|f_1 - 0.5| < 0.02$;*
4. *pass at every test of the NIST SP 800-22 and Dieharder batteries with χ^2-uniform p-value distribution.*

In addition, the analog hybrid-mixing stage of [22] §3.7 combines three independent physical sources (Barkhausen, Johnson–Nyquist, avalanche) so that successive blocks $\mathcal{E}^*_j$ are statistically independent of past blocks and of the messages being encrypted. Under these certified conditions, each block $\mathcal{E}^*_j \in \{0,1\}^\ell$ is modelled as drawn from the uniform distribution U_ℓ on $\{0,1\}^\ell$, in accordance with the hypothesis of [22], §2.5.

Theorem 6.2 (E* perfect secrecy under hardware certification)**.** *Let $\mathsf{CT} = \mathsf{TX} \oplus \mathcal{E}^*_j$ with $\mathcal{E}^*_j \in \{0,1\}^\ell$ ($\ell = |\mathsf{TX}|$) satisfying Lemma 6.1, and let $j \in \mathbb{N}$ be the entropy serial index. Then $\mathcal{E}^*_j$ acts as a one-time pad and the construction achieves Shannon perfect secrecy:*

$$\mathcal{I}(\mathsf{TX}; \mathsf{CT}) = 0, \tag{16}$$

and consequently, for any adversary $\mathcal{A}$ (classical or quantum, computationally unbounded) observing the public transcript (CT, j):

$$\Pr[\mathcal{A}(\mathsf{CT}, j) = \mathsf{TX}] = 2^{-\ell}. \tag{17}$$

Proof. The serial index j is an opaque routing label assigned at pre-assignment time, independently of TX: the mapping $j \mapsto \mathcal{E}^*_j$ is private to the recipient, so j carries no information about TX and $\mathsf{H}(\mathsf{TX} \mid j) = \mathsf{H}(\mathsf{TX})$. We therefore restrict the analysis to $\mathcal{A}(\mathsf{CT})$. By Lemma 6.1, $\mathcal{E}^*_j$ is uniformly distributed on $\{0,1\}^\ell$ and independent of TX; the hypotheses of Theorem 3.3 are satisfied, yielding $\mathcal{I}(\mathsf{TX}; \mathsf{CT}) = 0$ and (16). Since CT is then uniformly distributed regardless of TX, no function of CT correlates with TX and the optimal adversary strategy succeeds with probability exactly $2^{-\ell}$, giving (17). □

7 Entropy-Padded Transactions

7.1 The EPT Encryption Model

An *Entropy-Padded Transaction* is a tuple

$$\mathsf{EPT} = \big(\mathsf{CT},\ j,\ \mathsf{sig}\big), \tag{18}$$

where:

- $\mathsf{CT} = \mathsf{TX} \oplus \mathcal{E}^*_j$ is the one-time-pad-encrypted transaction, with $|\mathcal{E}^*_j| = |\mathsf{TX}|$.

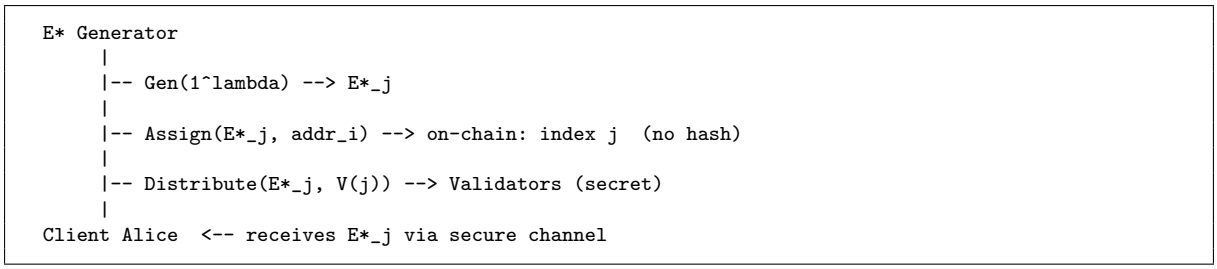

Figure 1: Entropy pre-assignment overview. Only the integer serial index j is recorded on-chain — no hash of $\mathcal{E}_j^*$ is ever published. The entropy $\mathcal{E}_j^*$ itself is distributed off-chain to the client and to the designated validator set $\mathcal{V}(j)$.

- $j \in \mathbb{N}$ is the *entropy serial index*: a sequential integer assigned by the $\mathcal{E}^*$ generator at pre-assignment time, used to look up the corresponding entropy block in the validator's local store. No hash function is involved; j is an opaque routing label.
- $\mathsf{sig} = \mathsf{ECDSA.Sign}(\mathsf{sk}_{\mathsf{Alice}}, \mathsf{CT} \| j)$ is the sender's standard Ethereum ECDSA signature.

Remark 7.1 (Padding)**.** *To prevent transaction-size side-channel leakage, all transactions are padded to a canonical length $L_{\max} = 2048$ bytes before encryption. Padding bytes are stripped by the validator after decryption.*

7.2 Entropy Pre-Assignment

Definition 7.2 (Entropy pre-assignment)**.** *An entropy pre-assignment scheme Π_{PA} consists of:*

1. $\mathsf{Gen}(1^\lambda) \to \mathcal{E}_j^*$*: a quantum entropy generator that produces a block $\mathcal{E}_j^* \in \{0,1\}^{L_{\max}}$ satisfying Lemma 6.1.*
2. $\mathsf{Assign}(\mathcal{E}_j^*, \mathsf{addr}_i) \to j$*: an off-chain procedure that assigns $\mathcal{E}_j^*$ to client address addr_i and records the pair (j, addr_i) on-chain, where j is a sequential integer index. No hash of $\mathcal{E}_j^*$ is published; the index j carries no information about the entropy value.*
3. $\mathsf{Distribute}(\mathcal{E}_j^*, \mathcal{V}(j))$*: a mechanism that delivers $\mathcal{E}_j^*$ to each validator $V_m \in \mathcal{V}(j)$—the designated validator set for entropy index j—through a confidential channel, such that validators store $\mathcal{E}_j^*$ locally without publishing it.*

Confidentiality requirement. The $\mathsf{Distribute}$ step must ensure $\mathcal{E}_j^*$ is not observable by any party outside $\{\mathsf{Alice}\} \cup \mathcal{V}(j)$. We state this as an open security requirement; the design of a concrete quantum-safe protocol satisfying it is left as future work.

One-time use invariant. Each entropy block $\mathcal{E}_j^*$ must be used for at most one transaction. The on-chain registry maintains a *spent* flag for every index j:

$$\mathsf{spent}[j] \leftarrow \mathsf{true} \quad \text{upon first inclusion of any EPT with index } j. \tag{19}$$

Any subsequent EPT referencing a spent index j is rejected by all honest validators.

Algorithm 1 Validator EPT verification

Require: $(\mathsf{CT}, j, \mathsf{sig})$, local entropy store $\mathcal{S}$
Ensure: Validated TX in mempool, or rejection

1: Verify $\mathsf{ECDSA.Verify}(\mathsf{vk_{Alice}}, \mathsf{CT}\|j, \mathsf{sig}) = 1$; else **reject**.
2: Verify $\mathsf{spent}[j] = \mathsf{false}$; else **reject**.
3: Look up $\mathcal{E}_j^* \leftarrow \mathcal{S}[j]$; if not found, **skip** (not assigned to this validator).
4: $\mathsf{TX}' \leftarrow \mathsf{CT} \oplus \mathcal{E}_j^*$.
5: Strip padding: $\mathsf{TX} \leftarrow \mathsf{TX}'[0 : |\mathsf{TX}'|_u]$ where $|\mathsf{TX}'|_u$ is encoded in the transaction header.
6: Run standard EVM transaction validation (nonce, balance, gas, signature); else **reject**.
7: Add TX to mempool.

7.3 Transaction Broadcast

Transaction submission proceeds as follows:

1. Alice retrieves her assigned entropy block $\mathcal{E}_j^*$ from her $\mathcal{E}^*$ dongle.
2. Alice pads TX to $L_{\max}$ bytes: $\mathsf{TX}' = \mathsf{TX}\|\mathbf{0}^*$.
3. Alice computes $\mathsf{CT} = \mathsf{TX}' \oplus \mathcal{E}_j^*$.
4. Alice constructs and signs:
$$\mathsf{EPT} = \big(\mathsf{CT},\ j,\ \mathsf{ECDSA.Sign}(\mathsf{sk_{Alice}}, \mathsf{CT}\|j)\big). \tag{20}$$
5. Alice broadcasts EPT to the peer-to-peer network.

7.4 Validator Decryption and Mempool Admission

Upon receiving EPT, each Validator Node $V_m \in \mathcal{V}(j)$ performs the following checks before admitting the transaction to its mempool:

7.5 Perfect Secrecy Guarantee

Theorem 7.3 (EPT perfect secrecy)**.** *Let* $\mathsf{EPT} = (\mathsf{CT}, j, \mathsf{sig})$ *be a valid Entropy-Padded Transaction. For any adversary* $\mathcal{A}$ *(computationally unbounded, including quantum) who observes* EPT *without access to* $\mathcal{E}_j^*$*:*

$$\mathcal{I}(\mathsf{TX}\,;\,\mathsf{CT}) = 0. \tag{21}$$

Moreover, $\mathcal{A}$ *has no advantage in guessing any bit of* TX *beyond random chance, including any bit of the ECDSA signature or public key embedded in* TX*.*

Proof. By Lemma 6.1, $\mathcal{E}_j^*$ is uniformly distributed on $\{0,1\}^\ell$ and independent of TX. The hypotheses of Theorem 3.3 are therefore satisfied for the construction $\mathsf{CT} = \mathsf{TX} \oplus \mathcal{E}_j^*$, yielding $\mathcal{I}(\mathsf{TX};\mathsf{CT}) = 0$ exactly. The index j is a sequential lookup label assigned independently of TX and carries no information about $\mathcal{E}_j^*$ or TX. The signature sig is computed over (CT, j) — neither field contains TX in plaintext. Therefore no function of $(\mathsf{CT}, j, \mathsf{sig})$ correlates with TX. In particular, the ECDSA public key $\mathsf{vk_{Alice}}$, which is part of TX, is information-theoretically hidden, so Shor's algorithm cannot be applied. □

Corollary 7.4. *Any network observer—including miners, full nodes, and relay nodes—who sees a broadcast* EPT *but is not in the designated validator set* $\mathcal{V}(j)$ *has zero information about the transaction content until* $\mathcal{E}_j^*$ *is released post-confirmation.*

7.6 Post-Confirmation Entropy Revelation

After a block containing a transaction with entropy index j achieves finality (as determined by the SyncroChain), the entropy $\mathcal{E}_j^*$ is published on-chain. This serves two purposes:

1. **Auditability.** Any node can now verify $\mathsf{CT} \oplus \mathcal{E}_j^* = \mathsf{TX}$ and cross-check the EVM state transition.
2. **Spent enforcement.** The on-chain record of $\mathcal{E}_j^*$ ensures it can never be reused as a one-time pad.

The revelation transaction is submitted by any member of $\mathcal{V}(j)$ within T_{reveal} blocks after finality; failure to reveal triggers a slashing condition on the responsible validators.

8 Proof of Quantum Work

8.1 Motivation

Classical Proof-of-Work (PoW) as used in Bitcoin requires miners to solve:

$$\mathsf{SHA\text{-}256}(B_h \| \mathsf{nonce}) < \mathcal{D}, \tag{22}$$

where B_h is the block header and $\mathcal{D}$ is the target difficulty. This protocol suffers two quantum-era pathologies:

1. Grover's algorithm gives any quantum miner a permanent $O(\sqrt{T})$ vs $O(T)$ advantage over classical miners, inevitably concentrating mining in quantum-equipped entities.
2. Bitcoin's energy consumption reached 175.87 TWh in 2024, an unsustainable environmental burden.

Proof of Stake (PoS) avoids the energy problem but concentrates power in wealthy validators. DAC adopts *Proof of Quantum Work (PoQW)* [8], which replaces SHA-256 mining with a genuinely quantum computational task that:

1. Cannot be efficiently performed on classical hardware.
2. Requires true quantum hardware (quantum annealers) to solve efficiently, creating an honest energy-security tradeoff anchored in quantum physics rather than classical hash power.
3. Reduces energy consumption by more than 99% compared to classical PoW.

8.2 Quantum Hash Function

PoQW replaces the classical deterministic hash with a *quantum hash function* whose output is defined by the probability distribution of a quantum annealer's measurement outcomes.

Definition 8.1 (Quantum hash function [8])**.** *Let $\mathcal{P}$ be a quantum annealing processor with N qubits. For block header data $x \in \{0,1\}^*$, define the quantum hash $\mathcal{H}_Q(x)$ as a length-n bit string derived from the expectation values of the final annealed state:*

$$\mathcal{H}_Q(x) = f\big(\langle\sigma_1\rangle_x, \langle\sigma_2\rangle_x, \ldots, \langle\sigma_n\rangle_x\big)\,, \tag{23}$$

where $\langle\sigma_i\rangle_x$ is the marginal expectation of qubit i in the steady state of the Ising Hamiltonian $H(x)$ derived from x, and f is a publicly specified deterministic post-processing function.

$\mathcal{H}_Q$ satisfies the following properties required of a PoW hash:

1. **Fixed output size:** $|\mathcal{H}_Q(x)| = n$ bits for all x.
2. **Avalanche effect:** flipping a single input bit causes $\approx n/2$ output bits to change (verified empirically in experiments [8]).
3. **Uniform distribution:** for random x, the output distribution of $\mathcal{H}_Q(x)$ is statistically close to uniform.
4. **Probabilistic (quantum property):** unlike a classical hash, two evaluations of $\mathcal{H}_Q(x)$ on the same input may yield different outputs, as each anneal samples from the same stationary distribution. The PoQW consensus protocol accounts for this (see below).
5. **Spoof resistance:** a classical computer cannot efficiently sample from $\mathcal{H}_Q$'s output distribution without simulating the full quantum annealing process, which is believed to require exponential classical time for the problem instances used.

8.3 D-Wave Quantum Annealing Implementation

The PoQW mining protocol proceeds as follows:

The proof π is verified by other nodes using the publicly specified f function, checking that the sample $\mathbf{s}$ is consistent with the stated Ising problem $H(B_h)$ and that $f(\mathbf{s}) < \mathcal{D}$. Full verification of quantum validity requires access to the same D-Wave quantum annealing hardware; light verification uses the classical f post-processing step and statistical consistency checks.

Infrastructure. DAC operates PoQW mining on four D-Wave Advantage processors distributed across North America, following the multi-region quantum annealing architecture described in [8]. This geographic distribution mitigates single-point-of-failure risks and ensures physical quantum hardware is genuinely exercised in the consensus process.

Algorithm 2 PoQW block mining

Require: Block header B_h (containing prev hash, Merkle root, timestamp, nonce), difficulty target $\mathcal{D}$

Ensure: Valid quantum proof π

repeat
  Increment nonce in B_h.
  Construct Ising Hamiltonian $H(B_h)$ from block header bits.
  Program D-Wave processor with $H(B_h)$.
  Anneal: obtain sample $\mathbf{s} \sim \Pr[\cdot \mid H(B_h)]$.
  Compute $h = \mathcal{H}_Q(B_h) = f(\mathbf{s})$.
until $h < \mathcal{D}$
Proof: $\pi = (\mathbf{s}, B_h, h)$.
return π

8.4 Viability Requirements

Following [8], we require that a PoQW protocol satisfy three formal viability conditions:

Definition 8.2 (Efficiency)**.** *A PoQW chain is* efficient *if the fraction of honest-miner blocks in the strongest valid chain exceeds a threshold $\phi > 1/2$ with high probability, even when a fraction $f < 1 - \phi$ of mining power is adversarial.*

Definition 8.3 (Immutability)**.** *A PoQW chain is* immutable *if, for any block B at depth d in the confirmed chain, the probability that B is removed is bounded by $e^{-\Omega(d)}$.*

Definition 8.4 (Security)**.** *A PoQW chain is* secure *if no adversary controlling strictly less than $1/2$ of the total quantum mining rate can successfully double-spend with probability greater than $\mathsf{negl}(\lambda)$.*

Theorem 8.5 (PoQW viability [8])**.** *Under the assumptions that (a) the quantum annealing process is spoof-resistant and (b) honest miners control a majority of quantum mining hardware, the PoQW protocol as implemented on a quantum annealer satisfies Definitions 8.2–8.4.*

The proof of Theorem 8.5 follows the security analysis in [8] and is based on a Markov chain model of block propagation. The key insight is that the probabilistic nature of quantum hashing does not compromise security: the protocol is designed to handle output variance from the annealer, and honest-majority assumptions mirror those of classical PoW security proofs.

8.5 Energy Efficiency

A current-generation quantum annealer consumes approximately 25 kW during operation. Bitcoin's entire network consumed 175.87 TWh in 2024, equivalent to an average power draw of $\approx 20\,\mathrm{GW}$. A PoQW network of n annealer nodes has power consumption:

$$P_{\mathrm{PoQW}} = n \times 25\,\mathrm{kW}. \tag{24}$$

For $n = 1000$ annealer nodes, $P_{\mathrm{PoQW}} = 25\,\mathrm{MW}$, a reduction of more than three orders of magnitude (99.9%) compared to Bitcoin PoW.

```
PHASE 0 -- Entropy Pre-Assignment (off-chain)
  E* Generator --> E*_j assigned to Alice  (serial index j)
  E* Generator --> E*_j distributed to V(j) validators
  On-chain:  record (j, Alice_addr)  -- no hash of E*_j

PHASE 1 -- Transaction Construction (Alice)
  TX  <-- construct EVM transaction
  TX' <-- pad TX to L_max = 2048 bytes
  CT  <-- TX' XOR E*_j           [OTP encryption]
  EPT <-- (CT, j, ECDSA.Sign(sk_Alice, CT||j))
  broadcast EPT to DAC P2P network

PHASE 2 -- Validator Decryption & Mempool
  V_m in V(j):
    verify ECDSA signature on (CT, j)
    verify spent[j] = false
    E*_j <-- local store[j]
    TX'  <-- CT XOR E*_j
    TX   <-- strip padding from TX'
    EVM validate TX (nonce, balance, gas)
    add TX to shard mempool

PHASE 3 -- Block Production (PoQW Mining)
  validator builds candidate block B_h
    (prev_hash, Merkle_root(mempool TXs),
     timestamp, nonce, entropy indices j...)
  Quantum annealer: H(B_h) --> anneal --> sample s
    h_Q = f(s) < D  --> valid proof
  broadcast block + proof pi = (s, B_h, h_Q)

PHASE 4 -- Finality (SyncroChain)
  MasterChain aggregates shard blocks
  SyncroChain seals epoch checkpoint
  spent[j] <-- true (replay protection)

PHASE 5 -- Entropy Revelation
  any V_m in V(j) publishes E*_j on-chain
  all nodes verify:  CT XOR E*_j = TX  (direct check, no hash)
  TX is now publicly auditable
```

Figure 2: Complete DAC transaction lifecycle.

9 Complete Transaction Lifecycle

Figure 2 illustrates the end-to-end flow of a transaction in DAC, integrating all three core components.

The lifecycle satisfies the following invariants at each phase:

- At the end of Phase 1: $\mathcal{I}(\mathsf{TX}; \mathsf{EPT}) = 0$ for any network observer outside $\mathcal{V}(j)$.
- At the end of Phase 2: TX is in the mempool only of $\mathcal{V}(j)$ nodes; no other party knows TX.
- At the end of Phase 3: block is valid under PoQW consensus.
- At the end of Phase 5: TX is fully public and auditable; $\mathcal{E}_j^*$ is permanently spent.

10 Security Analysis

10.1 Quantum Attack Surface

Shor's algorithm on ECDSA. DAC does not replace ECDSA. Instead, the OTP layer makes Shor's algorithm inapplicable during the broadcast window:

- By Theorem 7.3, the ECDSA public key $\mathsf{vk}_{\mathsf{Alice}}$ (contained in TX) satisfies $\mathcal{I}(\mathsf{vk}_{\mathsf{Alice}}; \mathsf{CT}) = 0$. A quantum adversary intercepting the broadcast has zero information about the public key and therefore cannot run Shor's ECDLP algorithm.

- After block confirmation, $\mathcal{E}_j^*$ is revealed and TX becomes public (Section 7.6). At this point the ECDSA public key is exposed. However, the transaction is already final and irreversible. Users following the recommended practice of generating a fresh address per transaction ensure that no future funds are at risk from post-confirmation key exposure.
- Transaction encryption uses an OTP ($\mathsf{CT} = \mathsf{TX} \oplus \mathcal{E}^*$), whose security is information-theoretic and entirely independent of any algebraic hardness assumption. No hash function is used in the signing or encryption path.

Grover's algorithm. Grover's algorithm is relevant to:

1. *Mining.* DAC replaces SHA-256 PoW with PoQW. PoQW is classically intractable by construction; the quantum annealing problem structure provides no Grover-exploitable search structure.
2. *OTP layer.* As proven in Corollary 3.4, Grover's algorithm provides no advantage against the OTP-encrypted ciphertext CT. No hash function is involved in the transaction encryption or signing path.

Quantum annealing attacks on ECDLP. The QUBO-based ECDLP attack of [7] requires an observable ECDSA public key. The OTP layer ensures that no public key is observable during the broadcast window. The attack is therefore inapplicable during the critical vulnerability period.

10.2 Classical Attack Resistance

51% / Majority mining attack. In PoQW, a 51% attack requires an adversary to control more than half of the total quantum annealing capacity in the DAC network. Unlike classical PoW (where mining hardware is commoditised), quantum annealers are scarce, expensive, and physically trackable. This makes majority mining attacks substantially more difficult than in classical PoW systems.

Front-running and MEV. By Theorem 7.3 and Corollary 7.4, the content of a broadcast EPT is invisible to validators outside $\mathcal{V}(j)$ until entropy revelation. MEV extraction (which requires knowledge of pending transaction content) is fundamentally disabled for parties outside the designated validator set. Within $\mathcal{V}(j)$, the threat remains; mitigating it via threshold secret sharing is identified as future work.

Replay attacks. Replay protection operates at two levels:

1. Standard EVM account nonce prevents replaying signed EVM transactions.
2. The on-chain *spent* flag on index j prevents reuse of any entropy block, ensuring each OTP is used exactly once.

Length leakage. As noted in Section 7.1, transactions are padded to $L_{\max} = 2048$ bytes before encryption, preventing transaction-size fingerprinting.

Table 4: Security comparison with existing post-quantum blockchains

Feature	**DAC**	**QRL**	**Ethereum+PQC**	**IOTA 2.0**
TX confidentiality	Perfect (OTP)	None	None	None
Signature scheme	ECDSA (OTP-shielded)	XMSS	Dilithium (proposed)	WOTS+
Hash dependency	None (signing layer)	SHA-256	SHAKE-256	SHA-3
Mining	PoQW (Annealer)	PoS	PoS	DAG+PoW
Entropy source	Physical QRNG	PRNG	PRNG	PRNG
EVM compatible	Yes	No	Yes	No
Anti-MEV	Yes (OTP)	No	No	No
Forward secrecy	Yes (OTP one-time)	No	No	No

10.3 Security Comparison with Existing Post-Quantum Blockchains

DAC is the only EVM-compatible blockchain offering *information-theoretic* transaction confidentiality during broadcast and mempool residence, with *no hash or lattice dependency* in the signing layer. The combination of physical quantum entropy, OTP encryption, ECDSA shielding, and quantum-native mining is, to our knowledge, unique in the blockchain literature.

11 Discussion

Security and independence of assumptions. DAC's core quantum resistance rests on a single security pillar: the *information-theoretic* OTP layer. No hash function, no lattice assumption, and no computational hardness claim is required for transaction confidentiality or signing. The OTP layer—grounded in physics rather than mathematics—remains unconditionally secure against any quantum adversary.

EVM compatibility trade-offs. Full EVM compatibility imposes constraints: the Ethereum ABI and RLP transaction encoding are preserved, meaning DAC inherits Ethereum's tooling ecosystem. However, the EPT broadcast format is incompatible with standard Ethereum mempools, requiring a modified P2P relay layer. The EVM execution layer runs on decrypted transactions after validator processing, ensuring that smart contracts execute on plaintext state as in standard Ethereum.

Liveness under entropy exhaustion. If a client exhausts their pre-assigned entropy pool without receiving a new allocation, they cannot submit transactions until the entropy assignment is refreshed. Entropy pool management (allocation frequency, revocation, and renewal) must be specified in the full protocol specification; this is outside the scope of this paper.

Adversarial validators. A validator in $\mathcal{V}(j)$ who learns $\mathcal{E}_j^*$ can read transaction TX. This is an accepted limitation of the current pre-assignment model. DAC's staking model (Table 3) creates economic disincentives for validator misbehaviour through slashing; a threshold-sharing mitigation is left as future work.

12 Conclusion

We have presented DAC (Dual Asset Chain), a quantum-resistant EVM-compatible Layer-1 blockchain addressing the three principal quantum threats to deployed blockchain networks: the algebraic attack (Shor's algorithm on ECDSA and RSA), the search-space attack (Grover's algorithm on PoW mining and hash pre-images), and the latent confidentiality attack (future quantum decryption of recorded broadcast transactions).

DAC's *Entropy-Padded Transaction* protocol achieves information-theoretic transaction confidentiality via a one-time pad sourced from physically random quantum-mechanical Barkhausen noise. We proved that this OTP layer satisfies Shannon perfect secrecy unconditionally (Theorem 7.3), making it immune to any quantum algorithm by the laws of information theory rather than computational assumptions.

DAC's *Proof of Quantum Work* protocol replaces classical SHA-256 mining with quantum hashing on quantum annealers, anchoring consensus security in genuine quantum computational difficulty while reducing energy consumption by more than 99% compared to Bitcoin.

Standard Ethereum ECDSA signing is preserved without modification. The OTP layer guarantees $\mathcal{I}(\mathsf{TX};\mathsf{CT}) = 0$, rendering Shor's algorithm inapplicable: a quantum adversary intercepting the broadcast observes only uniform random ciphertext and never has access to the ECDSA public key. No hash function and no new cryptographic assumption is required for the signing layer.

The primary open problem—the quantum-safe distribution of entropy to validator nodes—is stated as Definition 7.2 and left as a direction for future work.

DAC demonstrates that information-theoretic security, quantum-native consensus, and EVM compatibility can be achieved in a unified blockchain architecture. We believe this combination represents the most complete answer to the quantum threat currently described in the blockchain literature.

Acknowledgements

The authors thank the D-Wave Systems research group for their foundational work on Proof of Quantum Work and for making the PoQW white paper available as a reference.

Many thanks also to Massimo Morini and Vincenzo Vespri, for many suggestions.

References

[1] K. A. Jackson, C. A. Miller, and D. Wang, "Evaluating the security of CRYSTALS-Dilithium in the quantum random oracle model," *arXiv preprint arXiv:2312.16619*, 2023, quICS, University of Maryland / NIST. [Online]. Available: https://arxiv.org/abs/2312.16619

[2] Q. Yuan, M. Tibouchi, and M. Abe, "Quantum-access security of hash-based signature schemes," Cryptology ePrint Archive, Report 2023/556, 2023. [Online]. Available: https://eprint.iacr.org/2023/556

[3] D. McGrew, P. Kampanakis, S. Fluhrer, S.-L. Gazdag, D. Butin, and J. Buchmann, "State management for hash-based signatures," Cryptology ePrint Archive, Report 2016/357, 2016. [Online]. Available: https://eprint.iacr.org/2016/357

[4] T. Wiggers, K. Bashiri, S. Kölbl, J. Goodman, and S. Kousidis, "Hash-based signatures: State and backup management," IETF, Internet-Draft draft-wiggers-hbs-state-00, February 2024. [Online]. Available: https://datatracker.ietf.org/doc/draft-wiggers-hbs-state/00/

[5] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," *Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 124–134, 1994. [Online]. Available: https://doi.org/10.1109/SFCS.1994.365700

[6] J. Proos and C. Zalka, "Shor's discrete logarithm quantum algorithm for elliptic curves," *arXiv preprint quant-ph/0301141*, 2003. [Online]. Available: https://arxiv.org/abs/quant-ph/0301141

[7] Ł. Dzierzkowski, "The generalized method of solving the ECDLP using quantum annealing," *arXiv preprint arXiv:2410.08725*, 2024, faculty of Cybernetics, Military University of Technology, Warsaw. [Online]. Available: https://arxiv.org/abs/2410.08725

[8] M. H. Amin, J. Raymond, D. Kinn, G. Miller, and D-Wave Quantum Inc., "Blockchain with proof of quantum work," D-Wave Quantum Inc., Tech. Rep., February 2026. [Online]. Available: https://arxiv.org/abs/2503.14462

[9] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978. [Online]. Available: https://doi.org/10.1145/359340.359342

[10] S. Beauregard, "Circuit for shor's algorithm using $2n + 3$ qubits," *Quantum Information and Computation*, vol. 3, no. 2, pp. 175–185, 2003. [Online]. Available: https://arxiv.org/abs/quant-ph/0205095

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

[12] G. Brassard, P. Høyer, and A. Tapp, "Quantum cryptanalysis of hash and claw-free functions," *ACM SIGACT News*, vol. 28, no. 2, pp. 14–19, 1997, arXiv:quant-ph/9705002.

[13] C. E. Shannon, "Communication theory of secrecy systems," *Bell System Technical Journal*, vol. 28, no. 4, pp. 656–715, 1949. [Online]. Available: https://doi.org/10.1002/j.1538-7305.1949.tb00928.x

[14] T. Kadowaki and H. Nishimori, "Quantum annealing in the transverse Ising model," *Physical Review E*, vol. 58, no. 5, pp. 5355–5363, 1998. [Online]. Available: https://doi.org/10.1103/PhysRevE.58.5355

[15] National Institute of Standards and Technology, "FIPS 204: Module-lattice-based digital signature standard (ML-DSA)," NIST, Tech. Rep. FIPS 204, 2024. [Online]. Available: https://doi.org/10.6028/NIST.FIPS.204

[16] D. J. Bernstein and T. Lange, "Post-quantum cryptography," *Nature*, vol. 549, pp. 188–194, 2017. [Online]. Available: https://doi.org/10.1038/nature23461

[17] A. Hülsing, J. Rijneveld, and F. Song, "Mitigating multi-target attacks in hash-based signatures," in *Public-Key Cryptography – PKC 2016*, ser. Lecture Notes in Computer Science, vol. 9614. Springer, 2016, pp. 387–416. [Online]. Available: https://eprint.iacr.org/2015/1256

[18] P. Dartois, A. Leroux, D. Robert, and B. Wesolowski, "SQIsignHD: New dimensions in cryptography," in *Advances in Cryptology – EUROCRYPT 2024*, ser. Lecture Notes in Computer Science. Springer, 2024. [Online]. Available: https://eprint.iacr.org/2023/436

[19] L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé, "CRYSTALS-Dilithium: A lattice-based digital signature scheme," in *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2018, no. 1, 2018, pp. 238–268. [Online]. Available: https://doi.org/10.13154/tches.v2018.i1.238-268

[20] F. Liu, Z. Zheng, Z. Gong, K. Tian, Y. Zhang, Z. Hu, J. Li, and Q. Xu, "A survey on lattice-based digital signature," *Cybersecurity*, vol. 7, no. 7, 2024. [Online]. Available: https://doi.org/10.1186/s42400-023-00198-1

[21] Postquant Labs, "Quip.Network: Quantum unit interlock pathway — hash-based signature implementations (WOTS+ over Keccak-256)," Documentation and open-source libraries `hashsigs-rs`, `hashsigs-ts`, `hashsigs-solidity`, 2026. [Online]. Available: https://quip.gitbook.io/docs

[22] A. L. Gatta, "E*: Entropy-based shannon-level encryption via quantum noise sources," SQT - Swiss Quantum Technology, Tech. Rep., March 2025.

[23] National Institute of Standards and Technology, "SP 800-90B: Recommendation for the entropy sources used for random bit generation," NIST, Tech. Rep. SP 800-90B, 2018. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-90B