

Securing Elliptic Curve Cryptocurrencies against Quantum Vulnerabilities: Resource Estimates and Mitigations

Ryan Babbush,^{1,*} Adam Zalcman,^{1,†} Craig Gidney,^{1,‡} Michael Broughton,¹
Tanuj Khattar,¹ Hartmut Neven,¹ Thiago Bergamaschi,^{1,2} Justin Drake,³ and Dan Boneh⁴

¹*Google Quantum AI, Santa Barbara, CA 93111, United States*

²*Department of Computer Science, University of California Berkeley, Berkeley, CA 94720, United States*

³*Ethereum Foundation, Zeughausgasse 7a, 6300 Zug, Switzerland*

⁴*Department of Computer Science, Stanford University, Stanford, CA 94305, United States*

(Dated: March 30, 2026)

The expected emergence of cryptographically relevant quantum computers (CRQCs) will represent a singular discontinuity in the history of digital security, with wide ranging impacts. This whitepaper seeks to elucidate specific implications that the capabilities of developing quantum architectures have on blockchain vulnerabilities and potential mitigation strategies. First, we provide new resource estimates for breaking the 256-bit Elliptic Curve Discrete Logarithm Problem over the secp256k1 curve, the core of modern blockchain cryptography. We demonstrate that Shor’s algorithm for this problem can execute with either ≤ 1200 logical qubits and ≤ 90 million Toffoli gates or ≤ 1450 logical qubits and ≤ 70 million Toffoli gates. In the interest of responsible disclosure, we use a zero-knowledge proof to validate these results without disclosing attack vectors. On superconducting architectures with 10^{-3} physical error rates and planar connectivity, those circuits can execute in minutes using fewer than half a million physical qubits. We introduce a critical distinction between “fast-clock” (such as superconducting and photonic) and “slow-clock” (such as neutral atom and ion trap) architectures. Our analysis reveals that the first fast-clock CRQCs would enable “on-spend” attacks on public mempool transactions of some cryptocurrencies. We survey major cryptocurrency vulnerabilities through this lens, identifying systemic risks associated with advanced features in some blockchains such as smart contracts, Proof-of-Stake consensus, and Data Availability Sampling mechanism, as well as the enduring concern of “abandoned” assets. We argue that technical solutions would benefit from accompanying public policy and discuss various frameworks of “digital salvage” to regulate the recovery or destruction of dormant assets while preventing adversarial seizure. We also discuss implications for other digital assets and tokenization as well as challenges and successful examples of the ongoing transition to Post-Quantum Cryptography (PQC). Finally, we urge all vulnerable cryptocurrency communities to join the migration to PQC without delay.

* babbush@google.com

† viathor@google.com

‡ craiggidney@google.com

CONTENTS

I. Introduction	3
II. Quantum Attacks on the Elliptic Curve Discrete Logarithm Problem	5
A. Attack Types and Disclosure Models	5
B. Updated Quantum Resource Estimates	7
C. The Evolution of Offensive Quantum Capabilities	10
III. Attacks on Bitcoin’s Digital Signature Scheme	12
A. Public Key Exposure in Bitcoin Transactions	12
B. The Mechanics of a Quantum Attack on Bitcoin	13
C. The Origins of Quantum Vulnerabilities in Bitcoin	17
D. The Infeasibility of Quantum Attacks on Proof-of-Work	19
IV. Quantum Vulnerabilities of other ECDLP-Based Cryptographic Protocols	20
V. Quantum Vulnerabilities of the Ethereum Blockchain	22
A. Account Vulnerability	23
B. Admin Vulnerability	24
C. Code Vulnerability	26
D. Consensus Vulnerability	27
E. Data Availability Vulnerability	29
VI. Quantum Vulnerabilities of Other Blockchains	30
A. Bitcoin and Ethereum Derivatives	30
B. Privacy-Preserving Blockchains	31
C. Post-Quantum Blockchains	32
D. Stablecoins and Real-World Asset Tokenization	33
E. A Taxonomy of Quantum Risk Profiles for Distributed Ledgers	34
VII. Risks and Challenges in Migrating to Post-Quantum Cryptography	34
VIII. Dormant Digital Assets	36
A. The Challenge Posed by Dormant Assets	36
B. Bitcoin Community’s Options for the Challenge of Dormant Assets	37
1. Do Nothing, Burn, and Hourglass	37
2. Bad Sidechain	39
C. Public Policy Options for the Challenge of Dormant Assets	39
1. Regulated Destruction of Dormant Assets is Not Viable Policy	40
2. Digital Salvage	41
3. National Security Response	41
4. Engagement with Bitcoin Community	42
IX. Outlook	42
Acknowledgements	43
Financial Conflict of Interest Statement	43
References	43
A. Appendix: Zero-Knowledge Proof of Resource Costs	54
1. Low-Qubit Circuit Variant	54
2. Low-Gate Circuit Variant	54
3. Circuit Architecture and Approximate Correctness	55
4. Measurement Based Uncomputation and Kickmix Circuits	56
5. Verifiable Fuzz Testing via the Fiat-Shamir Heuristic	56
6. Cryptographic Attestation via SP1 and Groth16 SNARK	56
7. Code Availability for Zero-Knowledge Proof	57

I. INTRODUCTION

Quantum computers render tractable certain otherwise hard computational problems. This promises an immense benefit to many fields of science and technology [1–5], but also jeopardizes some widely deployed cryptosystems. Most importantly, cryptographically-relevant quantum computers (CRQCs) pose a threat [6–8] to widely deployed public-key cryptography [9] including the Rivest-Shamir-Adleman (RSA) cryptosystem [10] and Elliptic Curve Cryptography [11, 12]. The former arises from Shor’s efficient quantum algorithm for factoring integers [13] and the latter from Shor’s efficient quantum algorithm for the Elliptic Curve Discrete Logarithm Problem (ECDLP) [13]. Cryptosystems based on the ECDLP are used to establish authenticity and integrity of software, firmware and microcode updates [14, 15], to ensure authenticity and integrity of operating system images in secure boot [15], to secure web traffic using the Transport Layer Security (TLS) protocol [16], to ensure integrity and confidentiality in end-to-end encrypted messaging [17], to secure electronic passports and identification cards [18], to support passwordless and multi-factor authentication [19], to secure administrative access to cloud infrastructure using the Secure Shell (SSH) protocol [20], to authenticate domain name records [21], to implement cryptographic features in systems with low compute resources, such as embedded systems and Internet-of-Things devices [22], and in other critical applications [23].

Cryptocurrencies stand out among these applications of quantum-vulnerable cryptography for two reasons. First, in pursuit of efficiency and scale many blockchains depend crucially on ECDLP-based cryptography which uses keys that are almost an order of magnitude smaller than those of RSA at a similar security level [24] which means a smaller CRQC is required to break them. Second, unlike traditional finance which generally employs multiple safeguards, blockchains tend to offer no recourse against fraudulent transactions enabling unrecoverable theft with a forgery of a single digital signature.

Despite the critical dependence of cryptocurrencies on quantum-vulnerable ECDLP-based cryptography, the intersection of quantum computing and blockchain technology remains under-explored. Prior works have focused almost exclusively on Bitcoin [25–30] with only a few publications exploring the cryptocurrencies landscape more broadly [31–33]. Existing discussions rely on historical timelines for constructing real-world CRQCs and generally do not reflect state-of-the-art resource estimates for quantum algorithms. Furthermore, systemic vulnerabilities with implications for stablecoins and tokenization remain unexplored. This whitepaper seeks to paint a more comprehensive picture by sharing updated resource estimates for quantum attacks while widening the discussion beyond a focus on Bitcoin to other cryptocurrencies and modern developments, such as real-world asset tokenization. In our exploration of cryptographic innovations in modern cryptocurrencies, we come across novel quantum vulnerabilities and attack modes, such as the ability to use a CRQC to manufacture reusable classical exploits—a risk absent from Bitcoin.

We expect this document will be scrutinized by policymakers and stakeholders in the financial sector who will have differing levels of familiarity with these risks. However, we also intend for this to be a wake-up call to our colleagues in the quantum computing community. For many in the quantum community, this may be the first detailed exposure to the practical mechanics of blockchain vulnerabilities and the specific “onchain” consequences of the algorithms we study. Indeed, the quantum computing community maintains an odd relationship with cryptanalysis applications. It is seen as scientifically prestigious to develop and refine quantum cryptanalysis algorithms, yet it is often regarded as taboo to discuss the downstream implications of these algorithms or how they will influence the commercial quantum computing landscape. In this sense, the field has developed a culture of detachment from the real-world consequences of its work on breaking quantum-vulnerable cryptosystems.

This paper offers several technical and policy contributions that depart from conventional wisdom. We explain our position that progress in quantum computing has reached the stage where it is prudent to stop publishing details of improved quantum cryptanalysis to avoid misuse. Nevertheless, publishing trustworthy resource estimates remains essential to signal the proximity of quantum threats and motivate timely defenses. This tension between public transparency and the need to deny rogue actors specific attack details is a long-standing debate in computer security [34], leading to the “Responsible Disclosure” paradigm [35] where vulnerabilities are disclosed only after a remediation window. Adapting this practice to quantum cryptanalysis — where patching is exceptionally difficult and the threat timeline is shifting — we share updated resource estimates while withholding the specific quantum circuits. However, this approach gives rise to issues of trust and confidence.

The value of cryptocurrency assets is based on their digital security (assured by code and cryptography) and public confidence in the decentralized system (rather than in third party intermediaries as in traditional banking). The importance of public confidence in the system reflects the fact that cryptocurrencies, like traditional forms of money, are social institutions [36, 37] built on trust in the system and shared understanding of its rules and capabilities. While the digital security of cryptocurrencies can be attacked using CRQCs, public confidence can be undermined using Fear, Uncertainty and Doubt (FUD) techniques [38–40]. In this sense, unscientific and unsubstantiated resource estimates may themselves constitute a genuine or apparent attack on the system, which complicates responsible disclosure and technical discussion of quantum vulnerabilities in blockchain technologies. We address this communication challenge in two ways.

First, we reduce the FUD potential of our discussion by explicitly emphasizing and clarifying areas where cryptocurrencies are resilient against quantum attacks, sometimes despite widespread belief to the contrary, such as in the case of Bitcoin’s Proof-of-Work consensus mechanism’s immunity to quantum attacks using Grover’s algorithm, Zcash’s newest shielded pool’s resilience against quantum attacks on protocol parameters, and the widespread protection of public keys behind cryptographic hash functions on many blockchains, including Bitcoin. We also discuss successful PQC projects to highlight technical feasibility of continued operation of digital economy in a world with CRQCs.

Second, we rigorously substantiate our resource estimates by sharing a cryptographic zero-knowledge (ZK) proof [41, 42] that enables trustless third parties to cryptographically verify the estimates without access to the underlying attack details. Specifically, we publish a ZK proof that we have compiled two quantum circuits for solving the 256-bit ECDLP: one with 1200 logical qubits and 90 million Toffoli gates and one with 1450 logical qubits and 70 million Toffoli gates. ZK proofs have been previously suggested [43] as a tool to demonstrate the existence of security vulnerabilities (not related to quantum computers) without leaking the details necessary to carry out an attack. However, past work on this topic appears to focus on demonstrating the feasibility of ZK disclosures in principle, by applying techniques to previously reported vulnerabilities rather than using them to disclose a novel or improved attack as we do here.

On a standard superconducting architecture, using surface code error correction, we estimate these computations could be realized with fewer than half a million physical qubits (nearly a 20 fold reduction over prior estimates [44]). This reduction in quantum resources needed to solve ECDLP on a quantum computer reflects the general pattern of quantum algorithmic improvements, exemplified by the reduction in resources needed to break 2048-bit RSA [45, 46]. While we focus on compilation to a scaled up version of platforms that have been experimentally demonstrated, resource estimates could be reduced substantially by making more aggressive assumptions about hardware capabilities. Furthermore, our analysis gives the first clear indication that superconducting qubits could launch attacks within the average block time of Bitcoin and Bitcoin Cash, thus enabling “on-spend” attacks whereby a transaction is intercepted, the key is broken, and a fraudulent transaction is syndicated in the brief period of time before it is recorded on the blockchain. This prospect highlights the importance of migrating to Post-Quantum Cryptography (PQC) and of mitigation measures that thwart on-spend attacks, such as private mempools and commit-reveal schemes [47].

We delineate how superconducting, silicon, and photonic platforms, with their fast gates and short error correction cycle times, pose a distinct threat to active transactions that slower architectures like neutral atoms and ion traps likely cannot match. This architectural distinction informs our discussion on mitigation strategies. We explore two divergent scenarios: one where the first CRQCs are built on a superconducting, silicon, or photonic platform and the other where the first CRQCs are implemented in a neutral atom or ion trap architecture. In the latter case, we expect that the first CRQCs will only be able to implement “at-rest” attacks whereby the quantum attacker requires a substantial period of time to solve for the private key and is therefore only able to attack long-exposed public keys securing funds at rest on the blockchain. In this scenario, the first mitigation is not necessarily a full swap of the vulnerable cryptosystems, but rather a rigorous elimination of public key reuse and minimization of public key exposure. We note that while a distinction between attack types has been discussed in prior literature, sometimes under different names such as “short-range” and “long-range” attacks [25], we add to this discussion by associating attack types with specific quantum architectures. The ultimate path towards post-quantum security in blockchain technologies is technically clear, if logistically difficult: a full switch to PQC [48, 49]. We argue that steps towards this complex migration should begin immediately. In line with the defense-in-depth principle [50], we recommend that intermediate mitigation measures also be urgently adopted. Such measures are technically simpler than a full upgrade of the underlying cryptosystems which allows them to be deployed earlier.

The need for urgency arises from multiple considerations. Technical and logistical difficulties make migration to post-quantum cryptosystems a slow process. For some blockchains, this may be exacerbated by challenges involved in reaching sufficiently broad consensus. Key obstacles are related to resource costs of PQC which are substantially higher than those of ECDLP-based cryptosystems. In particular, in the Bitcoin community, some historical proposals that increased bandwidth requirements of running network nodes turned out to be controversial and led to disagreements and hard forks. Moreover, some of the current technical and financial trends in cryptocurrencies magnify quantum risks while exposing new funds and assets. Specifically, the quantum attack surface of blockchain-based systems continues to expand due to the introduction of new privacy and scalability features based on quantum-vulnerable cryptography. At the same time, financial developments, such as fiat-backed stablecoins and tokenization of other real-world assets (RWA), are projected to increase the pool of assets governed by smart contracts by nearly an order of magnitude by 2030 [51]. Most of this activity takes place on general-purpose blockchains for smart contracts, primarily Ethereum and to some extent Solana [52], with growing issuance on specialized blockchains, such as Algorand [53], Stellar [54, 55], and the XRP Ledger [56, 57] notable for protocol-level support for RWA tokenization [58–60]. The account model and smart contracts employed by these blockchains introduce new quantum vulnerabilities not present in Bitcoin and its derivatives.

Fortunately, there is a path to achieving post-quantum security for cryptocurrencies. PQC has become a mature cryptographic discipline: post-quantum cryptosystems have been proposed, scrutinized, implemented and deployed.

In fact, they are in active use protecting Internet traffic [61] and indeed securing blockchain transactions [62]. As we discuss in detail in later sections, some blockchains, such as the Quantum Resistant Ledger (QRL) [63, 64], Mochimo [65] and Abelian [66], rely exclusively on PQC. Others, such as Algorand, the XRP Ledger, and Solana, have made early experimental deployments of post-quantum protocols [62, 67].

However, forward-looking migration to PQC is not a panacea. Dormant digital assets, including those abandoned or inaccessible due to lost private keys, pose a distinct and critical challenge. We highlight the example of Bitcoin’s Pay-to-Public-Key (P2PK) locking scripts, which secure over 1.7 million BTC. The total amount of dormant quantum-vulnerable bitcoin may reach 2.3 million BTC when all script types are considered. Unlike active wallets that can migrate to new standards, dormant assets cannot be “fixed” via forks that enable PQC protocols for future transactions. They represent a fixed target — tens or hundreds of billions of dollars in value that will eventually become accessible to a quantum attacker. The community will soon face difficult, unprecedented decisions regarding the fate of these assets, forcing tradeoffs between the immutability of cryptographic property rights and the economic stability of the network.

The currently unclear legal status of the use of CRQCs to recover private keys of dormant P2PK assets creates significant risk that they would inevitably be seized by rogue actors. We consider policy responses to address this risk, including options for bringing the resulting financial gains into the formal, taxable economy and preventing them from falling exclusively to criminals or adversarial state actors. For example, governments may consider national security responses or may create a legal framework for “digital salvage”. This approach classifies the recovery of these assets as a regulated activity, analogous to recovery of sunken treasure. Ultimately, the fate of dormant quantum-vulnerable assets will depend on protocol changes (which may cause blockchains to fork), economic incentives (which will determine the relative value of forked assets) and government policy (which will constrain activities of regulated entities). The Bitcoin community retains control over protocol changes, but faces challenges related to the need for broad consensus. Therefore, it would be consistent for public policy to create a legal backstop by legalizing salvage while simultaneously recommending that the community exercise its autonomy to “burn” all salvageable coins.

The remainder of this paper is organized as follows. In [Section II](#), we present new quantum resource estimates for breaking the ECDLP together with a zero-knowledge proof that substantiates our claims while withholding technical details necessary to launch an attack. We also define the operational difference between “on-spend”, “at-rest” and “on-setup” attacks contrasting the capabilities of “fast-clock” quantum computers, such as those based on superconducting, silicon, and photonic qubits, on one hand, and “slow-clock” quantum computers, such as those based on neutral atoms and ion traps, on the other. In [Section III](#), we analyze the specific vulnerabilities of the Bitcoin blockchain, discussing the threat to digital signatures, explaining the detailed mechanics of “on-spend” and “at-rest” attacks, and the infeasibility of quantum attacks against its Proof-of-Work consensus mechanism. In [Section IV](#), we step beyond digital signatures and provide a high-level description of other ECDLP-based cryptographic protocols that enable new features that are a source of novel quantum vulnerabilities in modern cryptocurrencies. In [Section V](#), we examine the Ethereum blockchain, identifying distinct risks to its account model, smart contract governance, Proof-of-Stake validators, and Data Availability Sampling mechanism, highlighting how the emergence of a complex digital ecosystem, including Layer 2 solutions, stablecoins, and tokenization, substantially increases the attack surface for quantum threats. In [Section VI](#), we discuss smaller blockchains that are a source of intense financial and cryptographic innovation. We highlight novel quantum vulnerabilities they sometimes introduce as well as examples of successful deployment of PQC to protect cryptocurrencies and other digital assets. Then, in [Section VII](#), we describe risks and challenges associated with the migration to PQC. In [Section VIII](#), we shift from technical analysis to a discussion of public policy around quantum vulnerabilities, focusing on possible approaches to address the challenge of dormant assets. Finally, we conclude in [Section IX](#) with an outlook on the challenge and urgency of migrating to PQC.

II. QUANTUM ATTACKS ON THE ELLIPTIC CURVE DISCRETE LOGARITHM PROBLEM

A. Attack Types and Disclosure Models

At the heart of most cryptocurrency security lies the assumption that the ECDLP is hard to solve. In later sections, we will discuss specific blockchain functions that rely on this hardness assumption, including transaction validation in Bitcoin and Ethereum and the Proof-of-Stake consensus mechanism in Ethereum. Here, we outline briefly three types of quantum attacks on blockchain security and provide new bounds on the quantum resources required to execute such attacks. Based on these estimates, we discuss two scenarios of the emergence of CRQCs capable of exploiting cryptocurrency vulnerabilities. These scenarios will later anchor our discussion of possible mitigation strategies. In particular, the speed of attacks that early CRQCs can execute determines which blockchain vulnerabilities are exploitable and hence, informs the choice of intermediate mitigation measures.

We classify quantum attacks into three categories based on the necessary execution speed and on whether access to a CRQC is required for each attack or only once to prepare a universal reusable exploit:

1. **On-Spend Attacks:** Attacks targeting transactions in transit. When a blockchain user broadcasts a transaction, an attacker must derive the private key within the window of time allowed before the transaction is recorded on the blockchain. This requires a quantum computer fast enough to solve ECDLP within the transaction settlement time of the target blockchain which ranges from hundreds of milliseconds to a few minutes (e.g., about 400 milliseconds for Solana, about 12 seconds for Ethereum, about 10 minutes on average for Bitcoin). On-spend attacks are also known as “short-range” or “just-in-time” attacks [25, 68].
2. **At-Rest Attacks:** Attacks targeting public keys that remain exposed onchain or offchain for long periods of time, such as dormant wallets with reused keys. The attacker has days (or more) to derive the private key. At-rest attacks are also known as “long-range” or “long-exposure” attacks [25].
3. **On-Setup Attacks:** Attacks targeting fixed public protocol parameters that produce a universal reusable backdoor into a cryptographic protocol. The backdoor is created by means of a one-time off-line quantum computation on a CRQC and subsequent attacks utilizing it are executed on a classical computer. For example, an on-setup attack may involve the use of Shor’s algorithm to recover the so-called “toxic waste” discarded in a powers-of-tau trusted setup ceremony [69]. While the Bitcoin blockchain is immune to on-setup attacks, some scaling solutions, such as Ethereum’s Data Availability Sampling mechanism, and privacy protocols, such as Tornado Cash [70], are vulnerable to this especially insidious attack mode.

We note that there are other ways to classify quantum attacks on cryptocurrencies. In particular, one could categorize them by the security property they breach yielding a finer classification into attacks on asset ownership (e.g., transaction forgery in Bitcoin), consensus (e.g., deep blockchain reorganizations after compromising more than half of Ethereum validators), confidentiality (e.g., attacks on unlinkability of diversified addresses in Zcash), monetary integrity (e.g., creation of new coin in Mumblewimble or a vulnerable zk-rollup of Ethereum), solvency (e.g., peg collapse due to illegitimate Rootstock or stablecoin redemptions, stealthy drain of an anonymity pool in Tornado Cash), data availability (e.g., forging a data availability sampling proof in Ethereum), or governance (e.g., vote forgery in Cardano or in a Decentralized Autonomous Organization on Ethereum, asset takeover in a multisig bridge). In later sections, we discuss quantum attacks in terms of their onchain consequences, returning to all of the above examples. Here, we focus on the execution speed, because it has direct implications for the type of CRQC necessary to launch an attack.

Even though all universal quantum computer architectures are in principle capable of unlocking the same asymptotic quantum speedups for certain problems, such as the exponential speedup for ECDLP provided by Shor’s algorithm, the actual wall clock time they incur varies by orders of magnitude. These constant factors govern the ability of a given type of CRQC to launch on-spend attacks. The resource estimates we describe below indicate that superconducting [71], photonic [72, 73], and silicon spin qubit [74] CRQCs, with their fast gates and short quantum error correction cycles, will be able to solve ECDLP in the span of a few minutes and thus, to launch on-spend attacks. By contrast, the elementary operations on neutral atom and ion trap devices are about two to three orders of magnitude slower [75, 76]. As a consequence, we do not expect CRQCs in these slower architectures to be able to launch on-spend attacks. We will refer to the former as fast-clock CRQCs and to the latter as slow-clock CRQCs.

Traditionally, the gold standard for quantum resource estimation has been full transparency: publishing algorithmic innovations, logical circuits, and error-correction optimizations to ensure results are openly verifiable. Our team has historically adhered to this standard — for instance, in establishing the most efficient quantum algorithms and circuits for breaking 2048-bit RSA [45, 46]. However, the escalating risk that detailed cryptanalytic blueprints could be weaponized by adversarial actors necessitates a shift in disclosure practices. Accordingly, we believe it is now a matter of public responsibility to share refined resource estimates while withholding the precise mechanics of the underlying attacks. Transparency regarding the overall resource costs of quantum attacks is essential; if the community overestimates the resources required, the perceived “safety margin” may lead to a dangerous “wait-and-see” attitude. Given the technical trade-offs of PQC, such delays could make it difficult for vulnerable cryptocurrencies to execute an orderly transition.

This tension has sparked debate within the community. For example, Scott Aaronson, a prominent researcher and blogger in quantum computing, has recently oscillated between advocating for total non-disclosure [77] and full transparency [78]. We regard this debate as a quantum iteration of the older vulnerability disclosure controversy [34] which has roots in Kerckhoff’s Principle [79] and which, in modern times, sets the No Disclosure position against the Full Disclosure movement. In computer security, the debate has ultimately yielded a consensus known as the Responsible Disclosure or Coordinated Vulnerability Disclosure (CVD): a public disclosure preceded by an embargo to allow the vulnerable systems to be fixed. Variants of CVD with strict deadlines have been adopted by prominent computer security research institutions, such as CERT/CC at Carnegie Mellon University [80] and by Google’s Project

Zero [81], and were codified as an international standard ISO/IEC 29147:2018 [35]. While ordinary software embargoes typically last only a few months, publication is frequently deferred for hard-to-fix exposures, such as hardware flaws, operating systems bugs, or issues that require updates to technical standards [80]. Quantum vulnerabilities in ECDLP-based cryptography certainly qualify as hard-to-fix. Consequently, we believe it is consistent with the principles of responsible disclosure to withhold the quantum circuits necessary to launch quantum attacks.

Nevertheless, providing resource estimates without supporting data can appear unscientific, especially to the historically skeptical [82] and decentralized cryptocurrency community. We bridge this gap using a zero-knowledge (ZK) proof [41, 42]. By providing a ZK proof, we offer the community a rigorous way to verify the imminent threat to ECDLP without leaking sensitive details to potential quantum adversaries. We produce our ZK proof using a state-of-the-art software tool, the SP1 zero-knowledge virtual machine (zkVM) [83].

The statement that our ZK proof demonstrates is the following: we possess a classical reversible circuit of a specified size which on most inputs correctly computes point addition on the elliptic curve `secp256k1` [84, 85]. This is the primary bottleneck in Shor’s quantum algorithm and can be related to the cost of the overall algorithm in a straightforward fashion provided that one is using a “windowed arithmetic” technique for multiplying an elliptic curve point by a scalar [45]. Therefore, verifying circuits for this subroutine is sufficient to substantiate our resource estimates. Furthermore, for the purposes of solving the ECDLP on a quantum computer, it is sufficient for the elliptic curve point addition circuit to yield correct results on most, rather than all, inputs [86]. This framing does leak two pieces of information about our result: (i) that we are using a classical reversible circuit executed in quantum superposition for the elliptic curve point addition and (ii) that we are using windowed arithmetic. However, neither of these facts is surprising or unexpected. In fact, (i) has been true of all prior work on breaking ECDLP on a quantum computer and (ii) is a common element in all recent works [44, 87–90].

In order to produce the ZK proof using the SP1 zkVM, we implement a Rust program that takes a secret classical reversible circuit as input and checks that it correctly computes elliptic curve point addition on nine thousand random inputs. We commit to our secret quantum circuit via its SHA-256 cryptographic hash (a unique “digital fingerprint”). We generate the random test inputs using a cryptographically secure pseudo-random number generator (CSPRNG) — a SHAKE256 extendable-output function (XOF) initialized with the raw bytes of our circuit. This approach is closely related to the Fiat-Shamir heuristic [91], and can be shown to be sound under certain idealized assumptions about the hash function. More details are provided in the Appendix of this paper. The Rust code and the ZK proof are available in our Zenodo dataset [92].

B. Updated Quantum Resource Estimates

The landscape of published state-of-the-art compilations of Shor’s algorithm for n -bit ECDLP problem is complicated by the space-time tradeoff between the number of logical qubits (which governs how large the CRQC required for a quantum attack must be) and Toffoli gate count. Toffoli gates dominate the resource costs of fault-tolerant execution of Shor’s algorithm dictating its overall runtime [93, 94] and determining whether on-spends attacks are possible. The best published number of logical qubits is from a recent preprint by Chevignard et al. [87], where they estimate a logical qubit count of 1100 for $n = 256$ (and $3.12n + o(n)$ asymptotically). However, this small logical quantum memory footprint comes at the cost of more than 100 billion Toffoli gates. Such a large number of gates translates to both an exorbitant runtime, as well as a very large number of physical qubits (the number of physical qubits required depends on both the number of logical qubits and the number of gates).

On the other extreme, the best prior published Toffoli count is roughly 200 million by Litinski [44] (the title of the paper says 50 million not 200 million because it includes optimizations that reduce the amortized cost of solving multiple instances, whereas we are focusing on single instance attacks). Litinski’s lower Toffoli count comes at the cost of more than twice the number of logical qubits compared to Chevignard et al. Accounting for error correction overhead, the space required for Litinski’s approach is further estimated to be about 9 million physical qubits in a photonic architecture [44]. We compare these works and others to our own in Figure 1.

We are reporting here that our team has developed logical circuits to break ECDLP on elliptic curves over finite fields with n -bit prime modulus and n -bit group order requiring approximately $4.5n$ space. In Figure 2, we show the resources required for two variants of our approach (one optimized for qubit count, the other optimized for gate count), at different values of n . At $n = 256$ bits, the circuits use either 1200 logical qubits and 90 million Toffoli gates or 1450 logical qubits and 70 million Toffoli gates. In terms of the spacetime volume (a key resource which in particular drives the quantum error correction overhead), these estimates represent roughly an order of magnitude improvement over the most efficient prior work when applied to a single ECDLP instance. These are the resource estimates we demonstrate using the ZK proof discussed above. Our findings apply directly to ECDLP on `secp256k1` — an elliptic curve widely used in digital signatures on popular blockchains, such as Bitcoin and Ethereum.

Using standard assumptions about superconducting qubits hardware, such as planar degree-four connectivity and

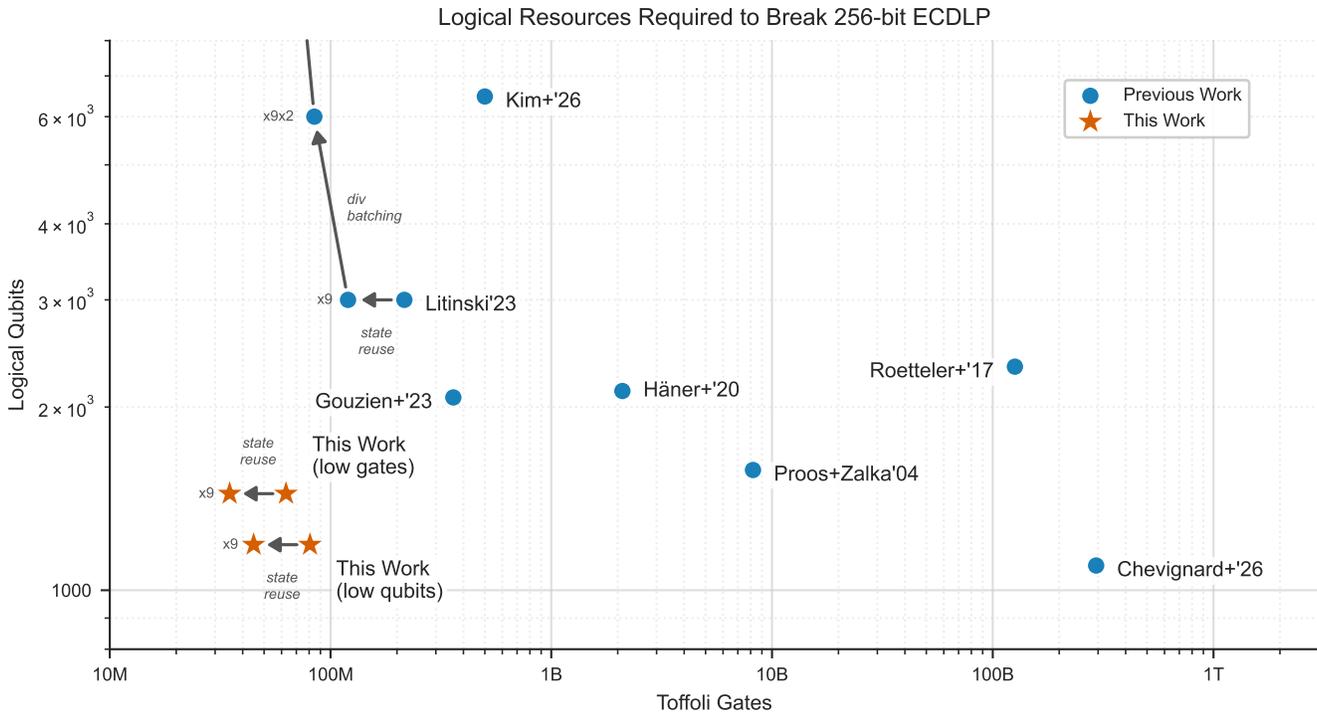


Figure 1. Comparison of logical quantum resources (number of logical qubits and Toffoli gates) required to break 256-bit ECDLP for the `secp256k1` curve, as reported by various prior works. The arrows illustrate the application of two algorithmic optimizations adopted from Litinski [44]: state reuse and div batching. State reuse refers to deriving multiple (here: nine) private keys from public keys in a single execution by reusing the initial phase estimation state. Div batching refers to running multiple (here: two for the visible point and nine for the point offscreen which corresponds to 2.7×10^4 logical qubits) instances of the algorithm in parallel and merging the modular division (inversion) operation across those instances; this is known as “Montgomery’s trick” in the cryptography literature [95]. Note that the main resource estimates for our approach quoted throughout this work apply to a single instance and do not include such optimizations. The plot represents resource estimates found in [44, 86–90, 96, 97].

10^{-3} physical error rates, running variants of the surface code architecture similar to the one assumed in [46], we calculate that these circuits can be executed on fewer than half a million physical qubits. This estimate is based on the reduced error correction overhead resulting from the smaller size of the optimized logical circuit and on the use of efficient storage of logical qubits in dense surface code configurations protected by “yokes” [98, 99]. In order to estimate the wall clock time of running our circuits, we assume they are executed in a reaction-limited fashion [94] — a style of quantum computation where execution speed is mainly limited by the reaction time of the control system. If we assume a control system reaction time of 10 microseconds (a standard figure for superconducting qubits [100]), and a 50% overhead per Toffoli gate, then 70 or 90 million Toffoli gates can be resolved in 18 or 23 minutes on a superconducting CRQC — close to the 10 minute average Bitcoin block time.

On-spend attacks, which affect all standard addresses in Bitcoin and similar cryptocurrencies, become even faster when we take into account the fact that a quantum computer can precompute the first half of the algorithm (which only depends on protocol parameters common to all addresses) and then wait in this “primed” state until a public key is available. This halves the time between a public key being revealed and its private key being resolved. Hence, we should estimate the time required to launch an on-spend attack starting from this primed state at the moment the public key is learned to be roughly either 9 minutes or 12 minutes. These facts imply that a superconducting CRQC capable of performing at-rest attacks against static holdings recorded on the blockchain would likely also be capable of executing on-spend attacks against active transactions. As we discuss in more detail later on, we do not expect meaningful scaling challenges between a quantum computer with 1200 logical qubits and one with 1450, so, in order to focus and simplify subsequent discussion, we assume that first-generation fast-clock CRQCs may be able to solve ECDLP on `secp256k1` and similar elliptic curves in about 9 minutes on average. Note that, if multiple primed machines are available, this duration can be further reduced. For example, according to Table 7.8 of Ref. [101], an attack with 11 primed machines would involve each machine performing 32 elliptic curve point additions rather than

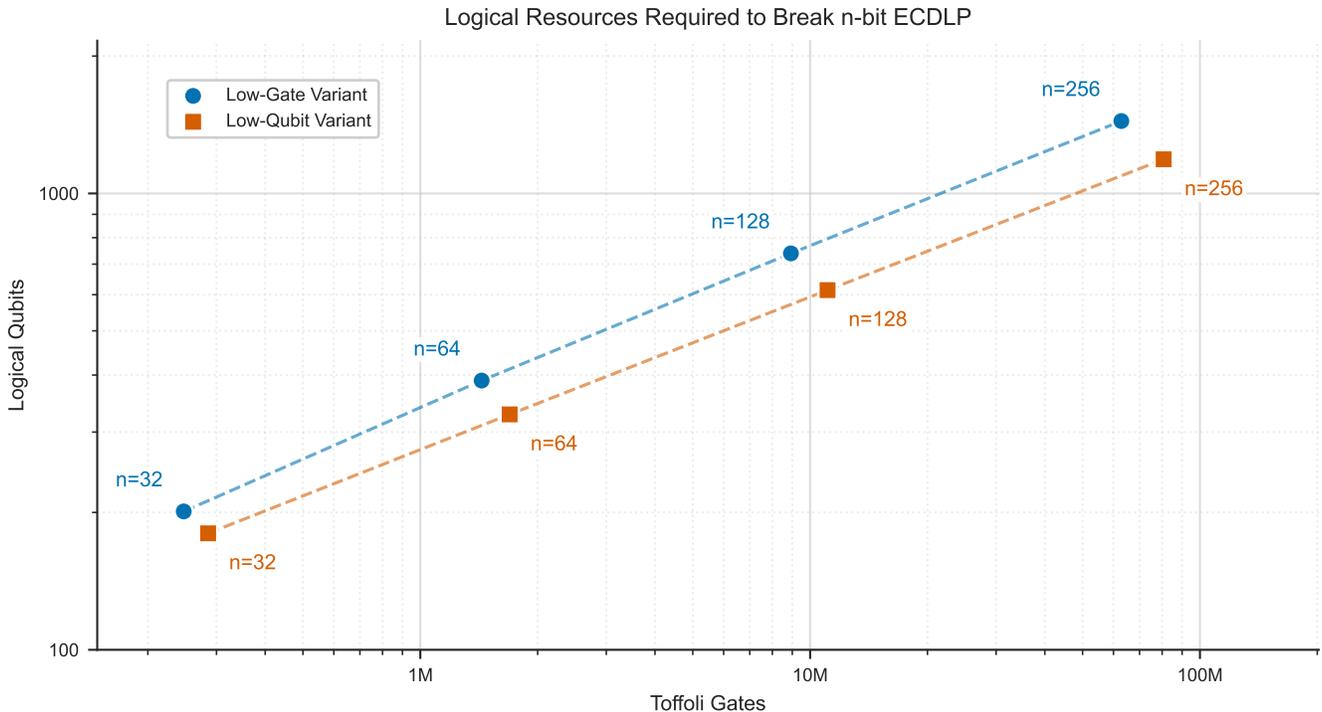


Figure 2. Logistical resources required to break n -bit ECDLP for curves with bit lengths of $n = 32, 64, 128, 256$.

one primed machine performing 208, a 6.5x speedup.

Executing 70 million Toffolis in 9 minutes would require the generation of half a million T states per second. A T state of sufficiently low error can be produced in 50,000 qubit-rounds [102]. In a fast-clock architecture with 1 microsecond error correction rounds, this implies allocating 25,000 physical qubits to magic production. This is a small qubit count compared to the half a million needed to run the algorithm at all, and this is why for fast-clock architectures we expect on-spend attacks to become possible at essentially the same time as at-rest attacks. By contrast, in a slow-clock architecture with 100 microseconds rounds or longer [75, 76], two and half million physical qubits would be needed for sufficient T state production. This qubit count is substantially larger than what is needed to run the algorithm slowly, so we currently expect that on slow-clock architectures at-rest attacks become possible before on-spend attacks. Nevertheless, it is plausible that future discoveries of new magic production techniques (e.g., new techniques for producing T states) that can be efficiently executed under the constraints of slow-clock architectures might enable them to launch on-spend attacks, too.

To further put these resource estimates into context we refer the reader to Figure 3 which first appeared in a recent perspective article from Google Quantum AI [103]. The lesson of Figure 3 is that advances in quantum algorithms and quantum error correction architectures continue to dramatically reduce the physical resources required to implement quantum applications (as the saying goes in cryptography: “attacks always get better” [104]). While resource estimates cannot drop indefinitely, the applications of focus in these plots (breaking RSA and simulating quantum chemistry) have been the focus of significantly more published research historically than quantum algorithms for breaking ECDLP so it may be the case that algorithms for those applications are closer to optimal than they are for ECDLP.

The high cost of fault-tolerant quantum computation stems largely from the overhead of quantum error correction, which demands a difficult balance of hardware-compatible connectivity, high noise tolerance, fast syndrome decoders, and expressive logical instruction set in order to provide a complete and feasible solution from an engineering perspective. While the well-understood surface code meets these requirements, it suffers from low efficiency — expressed by the logical-to-physical qubit ratio also known as the code’s encoding rate or simply rate — compared to other modern codes. For example, recent research into quantum Low-Density Parity-Check (qLDPC) codes [115, 116] offers an intriguing alternative by using long-range connections to increase the encoding rate. Although current resource estimates often default to the surface code for its proven decodability and the feasible demands it makes on hardware design, maturing qLDPC technologies could potentially further reduce physical qubit requirements [117, 118]. Recently, search for better quantum error correcting codes has been accelerated using AI techniques [119, 120].

The physical resource estimates we have discussed here (e.g., half a million physical qubits) assume relatively

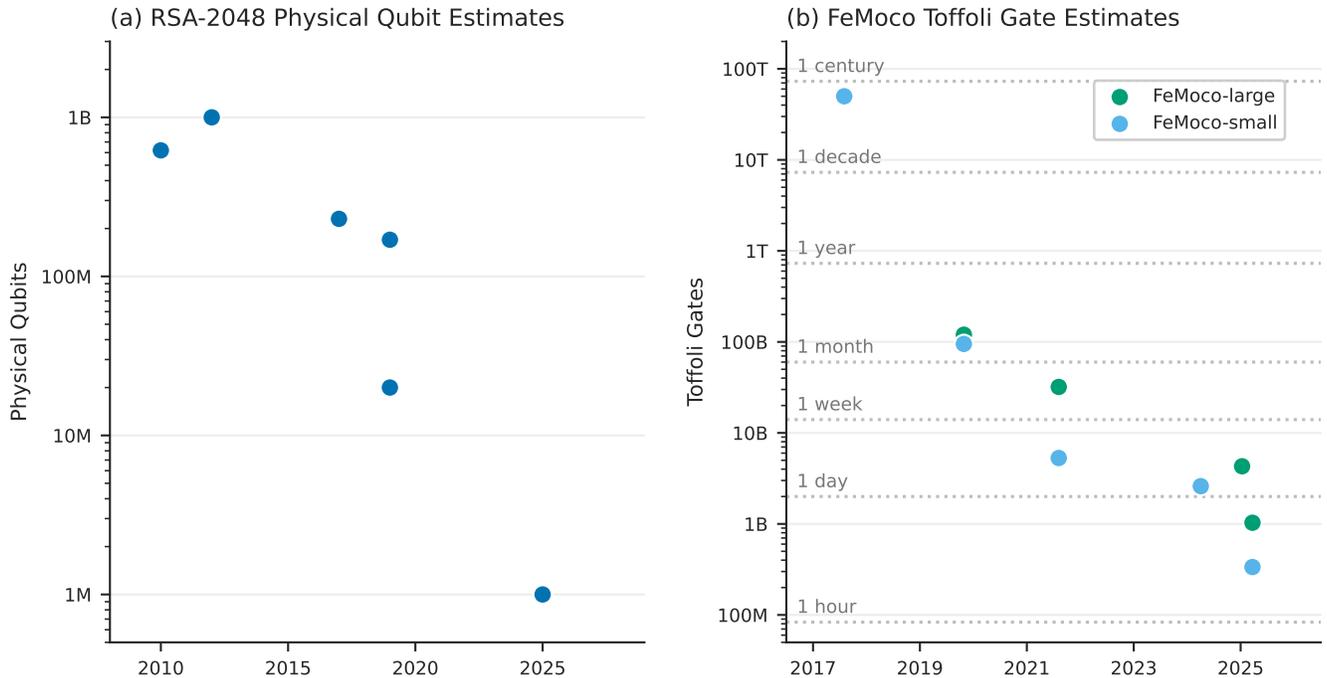


Figure 3. These figures, which first appeared in [3] illustrate that algorithms and error-correction research have dramatically decreased the resources required to solve important problems on quantum computers over the last decade. (Left) The number of physical qubits required of a superconducting qubit architecture running the surface code in the most advanced resource estimates for breaking 2048-bit RSA encryption, as a function of the year the manuscript was published. We note that some papers have claimed even fewer physical qubits by more aggressively changing hardware assumptions (e.g., by analyzing devices with higher connectivity or lower error rates) but here we elect to compare estimates that make comparable hardware assumptions. (b) Number of Toffoli gates required by the most advanced resource estimates for computing the ground state energy of the FeMoco molecule to chemical accuracy, as a function of the year the manuscript was published. Here, resource estimates are reported for the Reiher (“small”) [105] and Li (“large”) [106] FeMoco active space Hamiltonians. The left plot includes Refs. [8, 107–110] and the right plot includes Refs. [105, 111–114].

benign hardware capabilities, such as a planar architecture with degree-four connectivity and 10^{-3} physical gate error rates (i.e., consistent with a scaled up version of Google’s quantum processors that have been demonstrated experimentally [71]). More aggressive hardware assumptions — such as the “bicycle” architecture used for 2-gross qLDPC codes [115, 117] — could drop qubit counts closer to one hundred thousand physical qubits, but this approach requires non-local degree-seven connectivity that has yet to be demonstrated in actual superconducting qubit devices. Similarly, the recently introduced “Pinnacle” architecture [118] suggests breaking 2048-bit RSA with fewer than one hundred thousand superconducting qubits, but relies on even more challenging degree-ten, non-planar connectivity for codes that currently lack extensive simulations characterizing their performance, or efficient decoders, and would require nearly a month of runtime. The physical complexity of devices meeting these more demanding specifications makes them potentially more difficult to fabricate with sufficiently low physical error rates than larger devices with more benign engineering demands. Consequently, while a Pinnacle-style device would theoretically require even fewer than one-hundred thousand superconducting qubits to break 256-bit ECDLP, a planar degree-four device with half a million qubits remains the better understood and also likely more feasible engineering path. Nevertheless, the search for high rate codes compatible with engineering constraints is an active and promising research area [121].

C. The Evolution of Offensive Quantum Capabilities

Considering the significant speed difference between fast-clock and slow-clock quantum computing platforms, and because potential mitigation strategies depend on whether at-rest attacks become viable significantly earlier than on-spend attacks, we believe it would be prudent for the cryptocurrency community to develop contingency plans for two potential scenarios [122]:

- **Scenario 1:** On-spend attacks become viable about the same time as at-rest attacks. This is a possibility if the first CRQCs are fast-clock devices.
- **Scenario 2:** At-rest attacks become viable meaningfully sooner than on-spend attacks. This is a possibility if slow-clock quantum platforms advance more rapidly than fast-clock ones.

In part, the uncertainty concerning the two scenarios arises from the wide range of quantum computer architectures being actively developed across many research labs, tech companies and startups. For example, Google Quantum AI, IBM Quantum, Amazon, D-Wave, Rigetti and IQM are developing superconducting qubit architectures; PsiQuantum and Xanadu are building photonic quantum computers while Diraq and Intel are working on spin qubit devices. These three architectures feature fast elementary operations and are compatible with Scenario 1. Simultaneously, many companies, including IonQ, Quantinuum (a subsidiary of Honeywell) and Alpine Quantum Technologies are pursuing ion trap quantum processors while others, such as QuEra, Infleqtion, Atom Computing, Pasqal, and Logiqal are developing neutral atom devices. These two platforms have significantly slower elementary operations, so if the fast-clock architectures run into scaling challenges or if neutral atom or ion trap platforms scale rapidly, Scenario 2 could come to pass before Scenario 1.

It has been observed that technological change tends to occur in a two-stage evolutionary process in which an initial “era of ferment” characterized by high degree of innovation and technical variation eventually leads to the emergence of a “dominant design” and is followed by a period of gradual technological improvement [123]. For example, the car industry’s “era of ferment” (during the late 19th century) saw rapid innovation across a wide variety of cars powered by the steam engine, electric batteries and the internal combustion engine before the latter became the “dominant design”. In the stage of incremental improvement progress can be measured in terms of gradually increasing performance metrics, such as miles per gallon or the number of integrated circuit components in Moore’s law [124]. By contrast, quantum computing, with its broad variety of hardware platforms, is still in the “era of ferment” where simple models, such as counting physical qubits, fail to adequately capture technological progress. Instead, progress comes in discrete jumps corresponding to development of new internal capabilities and overcoming scaling challenges, e.g., by getting device error rates below the threshold for an error-correcting code [125] or implementing coherent interconnects for modular architectures [126–128]. Therefore, progress in quantum computing is better understood using a threshold model rather than in terms of the number of physical qubits.

Accordingly, other gradual measures of progress, such as the challenge ladder of increasingly difficult ECDLP instances ranging from 6-bit to 256-bit modulus and group order proposed in [90], also fail to adequately measure progress towards a CRQC and may fail to provide a reliable early warning. Indeed, if a leading quantum architecture encounters and overcomes all its scaling challenges before producing a device able to solve (for example) 32-bit ECDLP, then there may be little time between the breaking of 32-bit ECDLP and the breaking of 256-bit ECDLP. Furthermore, the community should not expect to see published demonstrations of the most advanced quantum error-correction architectures and quantum algorithms deployed to cryptanalytic problems. Thus, a successful public demonstration of Shor’s algorithm on a 32-bit elliptic curve should not be seen as a wake-up call to adopt PQC as much as a potential signal that PQC adoption has already failed.

Which of the two scenarios outlined above will come to pass depends on the scaling barriers affecting each platform and how quickly human ingenuity may overcome them. Given the uncertainty, we recommend that cryptocurrency communities urgently implement and deploy some of the relatively simple at-rest attack mitigations short of a full transition to PQC, such as removing vulnerable spending paths [129], updating wallets to warn against weak addresses, disabling support for Rootstock deposit transactions that expose public keys cross-chain, introducing rotation mechanisms for validator keys in Ethereum and staking and voting keys in Cardano [130]. Although such measures do not remove quantum vulnerabilities, they make attacks more difficult.

We have seen that while quantum hardware development efforts make progress on scaling quantum computers towards the “finish line” of cryptographical relevance, the work on quantum error correction and efficient compilation of quantum algorithms brings the finish line closer. We also argue that the finish line may become progressively more blurry and the current state of quantum capabilities increasingly more opaque as we approach CRQCs. Indeed, as progress in quantum computing lowers the barriers to entry, the final stages of the race to build a large fault-tolerant quantum computer may see “late-joiners” attempting a rapid breakout toward a CRQC, possibly accelerated by intellectual property theft or industrial espionage. Simultaneously, transparency of quantum computing research and development programs is likely to decrease as they get closer to large-scale commercially viable quantum computers. These factors may increase the uncertainty regarding the arrival time and the nature of the first CRQCs. Thus, it is conceivable that the existence of early CRQCs may first be detected on the blockchain rather than announced.

As we will argue throughout this piece, **the safest course of action for the cryptocurrency community is to begin preparing itself against quantum attacks immediately.** We will now turn to the discussion of the details of various quantum attacks against Bitcoin, Ethereum and other cryptocurrencies and how considerations of quantum architecture may come into play in mitigation strategies.

III. ATTACKS ON BITCOIN’S DIGITAL SIGNATURE SCHEME

A. Public Key Exposure in Bitcoin Transactions

The Bitcoin blockchain does not have a native concept of user accounts. Instead, it stores transactions, each of which is connected to other transactions via inputs and outputs. Units of bitcoin (BTC) exist as unspent transaction outputs (UTXOs), i.e., bitcoin that has been received but has not yet been spent, encumbered by a locking script, which ensures that it can only be spent by providing a valid digital signature matching the owner’s public key. When bitcoin is transacted from one wallet to another, the sender provides two pieces of cryptographic information: an unlocking script, which proves that they own the bitcoin being spent, and a locking script, which ties the bitcoin to the recipient’s wallet. One can think of the locking script as a cryptographic puzzle and of the unlocking script as its solution. Typically the puzzle involves digital signatures and is meant to be too hard for anyone to solve unless they know the private key.

Most common locking scripts are represented in a human-readable form as Bitcoin addresses. The overwhelming majority of bitcoin is locked using one of just seven standard types of scripts with various features, and degrees of vulnerability to quantum attacks [131–133]. A characteristic they all share is the use of ECDLP-based digital signatures to prove ownership of the bitcoin being spent. Transacting bitcoin requires one to present an unlocking script with a valid digital signature and to generate such a digital signature one needs to know the private key. Thus, to own bitcoin means to know the private key associated with the public key to which it is locked. CRQCs enable one to derive the private key from any public key, thereby upending Bitcoin’s current cryptographic model of ownership.

The varying degree of quantum vulnerability between various Bitcoin script types primarily arises from the differences in how they manage public keys. For example, P2PK locking scripts simply record the public key of the recipient directly on the blockchain. Later, when the recipient wishes to spend the bitcoin, they supply a digital signature in the unlocking script and the Bitcoin network validates that the signature was generated using the private key matching the public key from the locking script. Since the public key is disclosed directly in the locking script, it remains visible to everyone from the moment the bitcoin is received. Therefore, bitcoin locked by a P2PK script is vulnerable to both at-rest and on-spend attacks. A little over 1.7 million bitcoin (nearly 9% of all bitcoin) is secured by P2PK locking scripts.

In contrast, Pay-to-Public-Key-Hash (P2PKH) scripts are immune to most vectors of at-rest attack. Here, instead of disclosing the full public key in the locking script, the protocol only records its hash value (a unique “digital fingerprint”), on the blockchain. Later, when the bitcoin is spent, the unlocking script supplies both the public key and a digital signature. A CRQC can use Shor’s algorithm to derive the private key from a public key, but it cannot derive a public key from its hash value (that would require using the quantum computer with Grover’s algorithm to invert a hash function - see [Section III D](#) for more details on why this is not a credible threat). Therefore, bitcoin locked using a P2PKH script (and for which the public key is not leaked in other ways) is immune to at-rest attacks. However, spending such bitcoin does entail disclosure of the public key in the unlocking script, so it remains vulnerable to on-spend attacks.

Advanced users may require more complex spending conditions, e.g. m out of n signatures. The oldest standard script type supporting such transactions is Pay-to-Multisig (P2MS). Its quantum security properties are similar to P2PK, because it records the n public keys directly onchain in the locking script (and the unlocking script must supply m valid digital signatures corresponding to any m of these n keys). Here, a quantum attacker must derive m private keys rather than just one. However, since the public keys are exposed in the locking script, the attack can be launched offline. Pay-to-Script-Hash (P2SH) scripts potentially allow users to express even more complex spending conditions and hide them, together with any public keys they may contain, behind a hash. Therefore, bitcoin secured by P2SH locking scripts is immune to at-rest attacks, as long as the script is not reused or otherwise exposed before spending. In addition, for scripts that require a large number m of signatures, on-spend attacks by an attacker with a small number of CRQCs are made more difficult by the need to derive a large number of private keys in the ten minutes average block interval.

P2SH scripts and their modern replacements, such as P2WSH, can express spending conditions that do not depend on quantum-vulnerable cryptosystems. For example, in 2013, P2SH scripts were used to lock bounties for identifying weaknesses in hash functions used to compute hash values [134]. However, such scripts fail to cryptographically bind the unlocking script to the spending transaction, so any other Bitcoin user can copy the script to their own transaction before the original is recorded onchain. Thus, the effective spending condition such scripts express is: “anyone can spend once the puzzle is solved” rather than “whoever solved the puzzle can spend”. Therefore, the user must send the transaction to a trusted miner or mine it on their own.

Bitcoin’s Segregated Witness (SegWit) upgrade in 2017, introduced two new locking scripts: Pay-to-Witness-Public-Key-Hash (P2WPKH) and Pay-to-Witness-Script-Hash (P2WSH) which hide public keys and scripts behind hash values like the older P2PKH and P2SH scripts [135]. The new script types correspond to the modern Bitcoin

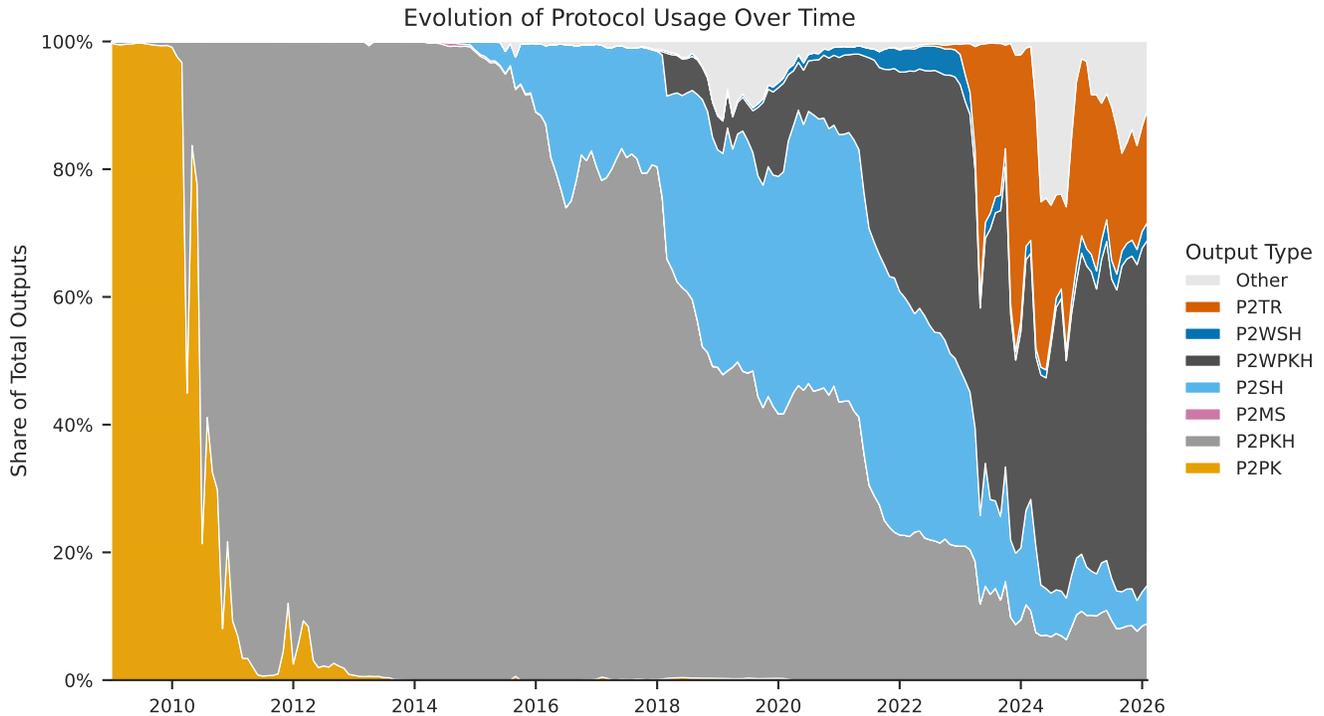


Figure 4. Evolution of Protocol Usage Over Time: The relative market share of transaction output scripts over time, highlighting the network’s migration through major protocol upgrades. The early “Satoshi Era” (2009–2010) is defined by the Pay-to-Public-Key (P2PK, orange), which was rapidly superseded by the industry-standard Pay-to-Public-Key-Hash (P2PKH, grey). The distinct inflection points in 2017 and 2021 correspond to the activation of the Segregated Witness and Taproot soft forks, respectively. The rapid expansion of P2WPKH (dark grey) and P2TR (red) demonstrates the ecosystem’s adoption of modern, weight-efficient cryptographic standards. Plot generated using data from `bigquery-public-data.crypto_bitcoin` [139].

addresses with the `bc1q` prefix. More recently, the Taproot soft fork in 2021 upgraded digital signatures from the old Elliptic Curve Digital Signature Algorithm (ECDSA) to the more flexible and efficient Schnorr signatures (both of which are based on ECDLP and thus vulnerable to quantum attacks) and introduced a new script type called Pay-to-Taproot (P2TR). This new versatile script type corresponds to Bitcoin addresses with the `bc1p` prefix and uses tweaked keys to simultaneously support two spending mechanisms called “key path spend” and “script path spend” which can provide many alternative complex spending conditions without forcing the user to reveal any of those conditions except the one they actually use at spend time. Thus, P2TR enhances privacy and supports advanced contracts and even promises to enable onchain verification of Turing-complete computations performed offchain [136]. However, P2TR stores the tweaked public key directly in the locking script without hiding it behind a hash value which brings back the quantum vulnerability of P2PK and P2MS addresses and therefore, from the standpoint of quantum security, constitutes a security regression [30, 137].

A recent draft Bitcoin Improvement Proposal BIP-360 [129] puts forward a new script type called Pay-to-Merkle-Root (P2MR) which is essentially P2TR with the quantum-vulnerable key path spend removed. This offers many benefits of P2TR without re-introducing vulnerability to at-rest attacks. P2MR, like all other standard script types in Bitcoin, is currently vulnerable to on-spend attacks. Thus, at present, P2MR constitutes a security patch for the Taproot regression. In future, when post-quantum opcodes are introduced in Tapscript [138], it may become fully quantum-secure. See Figure 4 for the relative prevalence of the standard Bitcoin locking scripts over time and Figure 5 for the total bitcoin locked by each type of script.

B. The Mechanics of a Quantum Attack on Bitcoin

An at-rest attack on a Bitcoin address that exposes the public key, such as P2PK or P2TR, proceeds as follows. First, the attacker reads the public key from any past transaction on the blockchain in which the address received bitcoin. Next, the attacker uses their CRQC to derive the private key from the public key. That private key is then

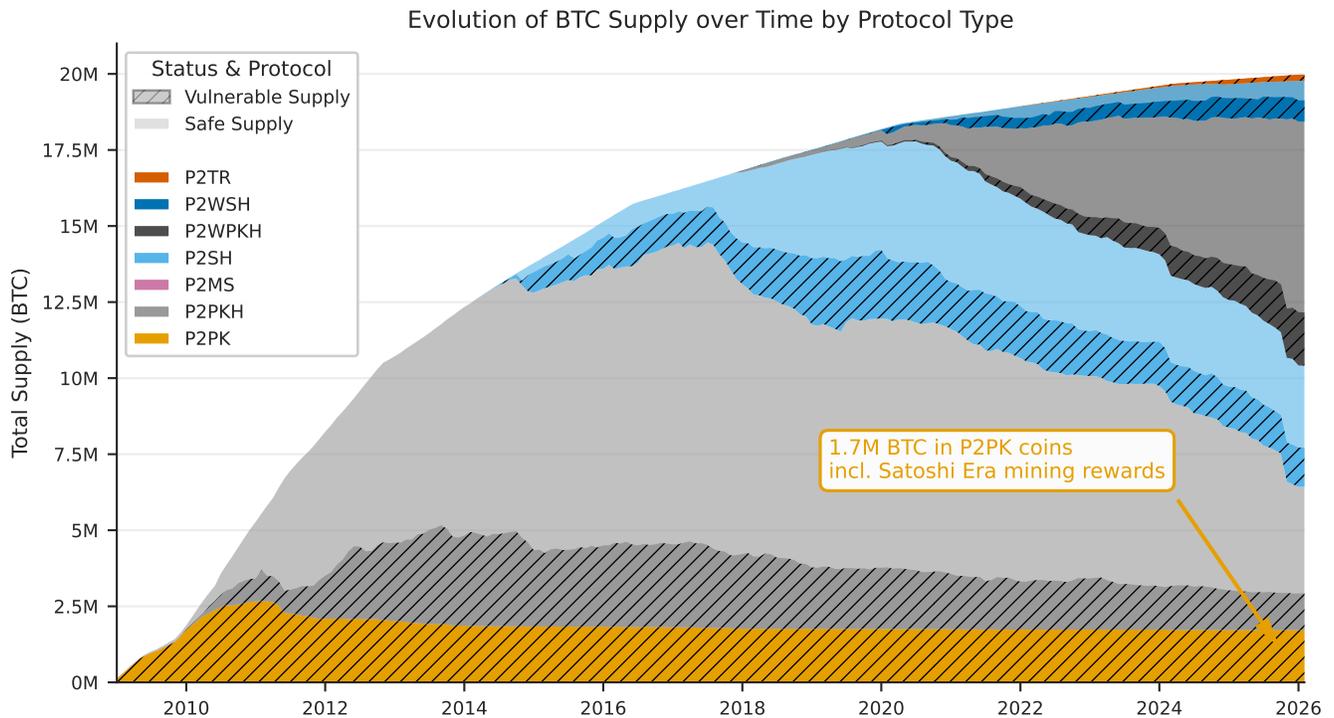


Figure 5. Evolution of BTC supply over time by protocol type. Quantum vulnerable balances are shown in shaded regions for each protocol. P2PK, P2TR and P2MS are considered 100% vulnerable. The remaining script types are considered vulnerable from having re-used keys if (Addresses that have appeared in an Input) AND (Currently hold a balance in Unspent Outputs). In the case of P2SH and P2WSH, we make the simplifying assumption that if one compromises the script, that it will ultimately equate to being able to steal the bitcoin (in some small number of cases this may not be true). At the time of writing ~6.9M total bitcoin across all protocols are vulnerable. Plot generated using data from bigquery-public-data.crypto.bitcoin [139].

used to authenticate a forged transaction locking the bitcoin behind a UTXO that the attacker controls. An at-rest attack on an address that hides the public key behind a hash value, such as P2PKH or P2WPKH, proceeds slightly differently. First, the attacker reads the public key from any past transaction on the blockchain in which the address spent bitcoin. Next, as before, they use their CRQC to derive the private key from the public key. Thus, absence of any spend transactions (and other public records disclosing the public key) denies the attacker the knowledge of the public key, thwarting the at-rest attack.

To understand the on-spend attacks, it is necessary to provide more background on how Bitcoin transactions are executed. First, the spender typically broadcasts their digitally signed transaction to the Bitcoin network where it is placed in a queue called the “public mempool” until a miner bundles it together with other transactions into a unit called a “block”. In order to have the block accepted by the network (and earn the block reward), the miner must solve a cryptographic puzzle - specifically, they must find a number that when included in this block causes its hash to have a value less than a certain difficulty threshold. This takes a considerable amount of computational work. Once the miner has found a solution to the puzzle, they then broadcast the completed block to the network whose nodes check that the block is valid (i.e., correctly formatted, consistent with prior blocks, includes only authentic transaction signatures, contains a correct solution to the cryptographic puzzle, etc.). When other miners construct their own blocks, they choose an “ancestor” block to build on top of, and the block reward incentivizes them to choose the tip of the longest valid chain of blocks. As a result, Bitcoin transactions are recorded in a ledger taking the form of a chain of blocks (hence, “the blockchain”). Due to inevitable network delays and compute power disparities, at any one time, miners might disagree on the last few blocks but their desire to arrest the growth of sunk costs causes them to eventually switch to the longest chain. The more blocks a transaction is buried under, the harder it is to reverse it with the cost of reversal growing exponentially in the number of blocks [140]. In practice, a transaction recorded a few blocks away from the tip of the blockchain is regarded as “finalized”. The longest chain rule provides a consensus mechanism while the requirement that blocks include Proofs of Work achieves Sybil resistance, i.e., a means of preventing a single malicious entity from easily gaining disproportionate control or influence by creating a large number of nodes. However, following informal convention we refer to the combination of the longest-chain consensus

rule with the Proof-of-Work Sybil resistance as the Proof-of-Work consensus mechanism.

In an on-spend attack, the quantum attacker uses the public mempool as the source of public keys. By the time a transaction is in the mempool, the public key associated with the address it is locked to must be visible, either on the blockchain (as is the case for P2PK and P2TR coins) or in the mempool (as is the case for P2PKH and P2WPKH coins) so that the network can check that the transaction is valid. For example, for a transaction spending P2PKH coins, the public key must be broadcast so that Bitcoin nodes can confirm that the public key indeed hashes to the value that is recorded on the blockchain in order to authenticate the digital signature. Critical for this discussion is the fact that it takes a nontrivial amount of time for a given transaction to be added to a block. On average, a new block is mined about every ten minutes. However, the mining procedure is a stochastic Poisson process, giving rise to an exponential distribution for the time between blocks mined; accordingly, the standard deviation is also about ten minutes. The standard deviation of the mempool wait time is even higher; at times of high network congestion, transactions can take days to process [141]. During that time, a quantum attacker can extract the public key from the mempool, use Shor’s algorithm to solve for the private key, and then use that private key to broadcast an alternative signed transaction attempting to move the Bitcoin to the attacker’s wallet. This “forged” transaction could be picked up by miners and added to a block that is ultimately finalized before the original transaction is recorded on the blockchain, thus successfully stealing the bitcoin.

A user spending bitcoin can increase a transaction fee in order to incentivize miners to prioritize their transaction when they construct a new block. Transaction fees are quite important at times of high congestion on the blockchain and, together with a small amount of newly minted bitcoin called “block subsidy”, form part of the mining reward. The relative importance of transaction fees in mining rewards gradually increases, because the block subsidy shrinks as the system slowly approaches the global monetary supply cap of 20,999,999.9769 BTC [142]. An attacker can flood the mempool with high-fee transactions to gain time for their CRQC to derive the user’s private key if needed.

In addition, miners can derive a supplemental revenue stream, known as Miner Extractable Value (MEV) [143, 144], from their ability to reorder, insert and remove transactions as they build new blocks. In particular, miners can sometimes extract profit from their control over which transactions are included in the next block, for example by front-running a large trade or benefitting from arbitrage. While MEV is more prevalent on blockchains with more complex financial ecosystems, such as Ethereum, the presence of quantum attackers may provide new MEV opportunities for miners. Bitcoin users might attempt to protect against on-spend attacks by offering miners a high fee to make sure that their transaction has minimal latency. However, a sufficiently fast quantum attacker could choose to make their fee even higher to increase the odds that their transaction is finalized ahead of the original. Such bidding wars might add a new component to the MEV revenue stream.

The timing and success probability of on-spend attacks are depicted in Figure 6. The probability of success is slightly less than 41% under the idealized assumptions stated in the caption. However, if our current estimate of around 9 minutes per private key derivation on a superconducting qubit CRQC could be brought down by a few minutes, then even high-fee transactions at times of low congestion would be likely to get compromised by a quantum attack. As discussed earlier, it is likely that the quantum attack time will indeed come down further, first due to further theory advances reducing the cost of Shor’s algorithm, and then due to increases in the number of qubits available to CRQCs relative to our present assumptions (our current quantum algorithm trades off time in order to fit on as few qubits as possible; doubling the number of qubits available to it would double its speed). The cost of quantum attacks on multi-signature transactions, such as P2WSH, increases linearly with the number of signatures required. However, these attacks parallelize trivially, so an attacker with a sufficient number of CRQCs can achieve the probability shown in Figure 6. Reducing block interval makes on-spend attacks significantly harder. For example, we estimate the success probability of an on-spend attack using a superconducting qubit CRQC against Litecoin, a Bitcoin derivative which targets 2.5 minutes between blocks, to be less than 3%. The chance of a successful on-spend attack on Zcash, whose target block time is 75 seconds, is less than one in thirteen hundred and on Dogecoin, whose target block time is 1 minute, is less than one in eight thousand.

Once the private key is derived, the attacker can engage the legitimate owner in a ‘Replace-By-Fee’ (RBF) bidding war. Because the attacker is stealing funds they do not own, they can rationally bid transaction fees up to incentivize miners to prioritize their theft. The victim might be forced into a ‘scorched earth’ scenario: to outbid the attacker and prevent the theft, they may have to offer a fee that may consume virtually the entire value of the asset. In this state, the transfer of wealth shifts from the victim to miners, rendering the asset effectively lost even if the theft is technically thwarted. In reality, mining operations are real-world businesses that can set up their own policies for handling such unusual situations. This creates an opportunity to seek outcomes that are more favorable to Bitcoin users than the simple scorched earth scenario suggests. In any case, if a quantum attacker cannot break the cryptography in less than say, 30 minutes or an hour, it seems likely that with a sufficiently high fee (or when the network congestion is sufficiently low), any transaction can be protected from a CRQC.

Legacy Bitcoin software allowed spenders to disable RBF for their transactions and implemented a relay policy known as the “first-seen” rule. RBF was also disabled in Bitcoin Cash following its hard fork from Bitcoin. Under

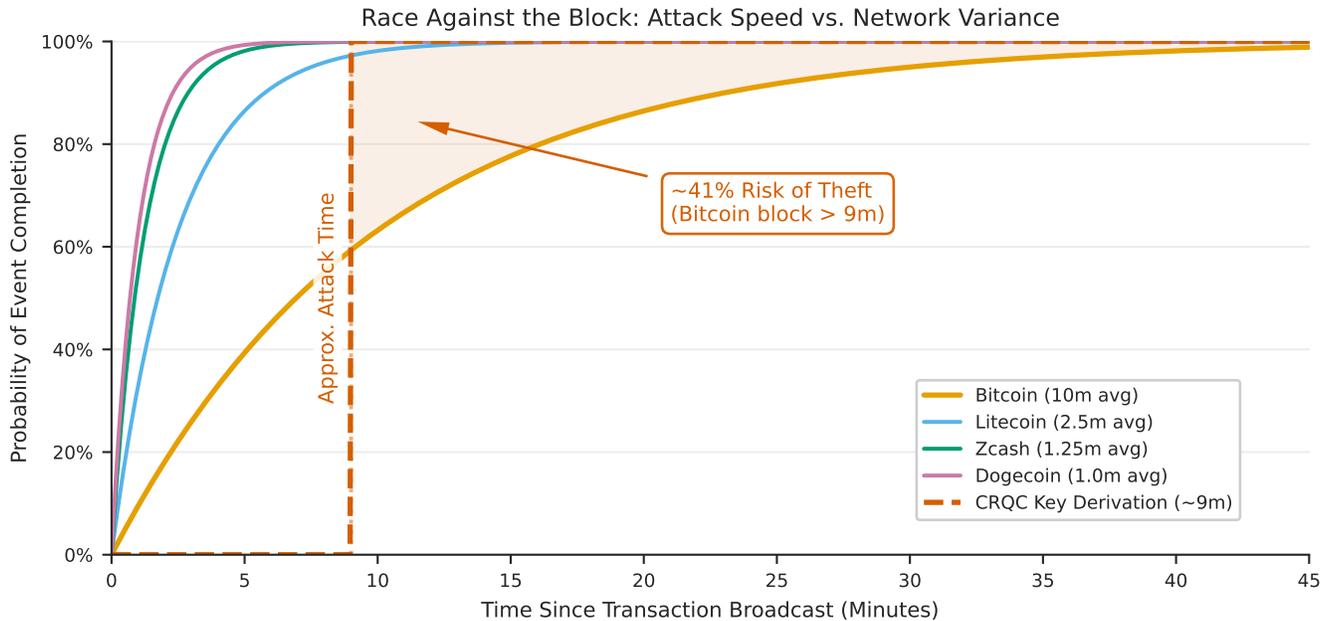


Figure 6. Risk that an on-spend quantum attack using a superconducting qubit CRQC taking approximately 9 minutes to derive a private key succeeds against Bitcoin (with average block time of 10 minutes), Litecoin (2.5 minutes), Zcash (75 seconds), and Dogecoin (1 minute). This plot assumes that the target transaction requires one digital signature (e.g., P2WPKH) and the public key is syndicated to the attacker almost immediately and that once broken (approximately at the ninth minute), the attacker can successfully insert a “forged” transaction by offering a higher fee to miners, displacing the original transaction (these assumptions are favorable to the attacker). But we are also assuming zero network congestion, which is favorable to the legitimate cryptocurrency spender. In practice, an attacker can create congestion artificially by flooding the mempool with high-fee transactions.

the first-seen rule, nodes encountering conflicting transactions would only relay the first one they saw rather than the one with the highest fee. However, the first-seen rule is sensitive to propagation latency and, more importantly, exists outside the consensus rules and is therefore unenforceable by the network. Indeed, miners have a financial incentive to prioritize high-fee transactions over first-seen transactions regardless of whether RBF is enabled or disabled. Consequently, disabling RBF was never a reliable protection against on-spend attacks.

Some Bitcoin mining operations allow users to bypass the public mempool by submitting transactions directly to miners for a fee [145]. Currently, these services are used to enhance privacy and efficiency and to enable users to utilize transactions that adhere to consensus rules, but are non-standard (such transactions can be included in valid blocks but are normally not propagated by the network). If Bitcoin fails to migrate to PQC before on-spend attacks become viable, these services could supplant the public mempool as a trust-based mechanism for achieving a moderate degree of quantum safety. However, like other mitigation measures against on-spend attacks, such as high fees and multi-signature scripts, these services reduce the quantum risk without eliminating it completely. For example, in the event of a blockchain reorganization, the orphaned block produced by the user’s trusted miner becomes visible on the network potentially exposing the user’s public key to a quantum attacker.

On-spend attacks may also interfere with Bitcoin consensus. The first time that the public key of a high value Bitcoin address appears in a block on chain, a quantum attacker could offer miners a reward to not mine on this block, thereby causing the block to become orphaned. This effectively cancels all transactions in this block. In parallel, the attacker would use its CRQC to break the just published public key included in the orphaned block. Once it has the secret key it issues a transaction to steal the assets from the target address, and this transaction will become part of the next (non-orphaned) block, thereby transferring the assets to the attacker. This reorganization of the chain is caused by the attacker, funded by the stolen funds, and is made possible by the CRQC.

C. The Origins of Quantum Vulnerabilities in Bitcoin

All quantum attacks on Bitcoin transactions are ultimately based on the ability of a CRQC to derive the private key from a public key. The most important distinction between different types of quantum vulnerabilities in Bitcoin is how the public keys become known to an attacker. Broadly speaking, there are four ways the attacker can learn the public key: they can retrieve it from a locking script or an unlocking script recorded on the blockchain, from a pending transaction in the mempool, or from offchain records. These possibilities correspond to the following vulnerabilities:

1. **Weak Address** vulnerability arises when the locking script directly reveals the public key of the wallet that is receiving bitcoin by recording it on the blockchain. In this case, the attacker has a long period of time to try to solve for the private key. The locked bitcoin remains vulnerable until it is spent. Weak Address vulnerabilities are caused by the use of locking scripts that leave public keys unprotected, namely P2PK, P2MS (in case of old assets) and P2TR (in recent transactions).
2. **Address Reuse** vulnerability arises when the unlocking script recorded in a spend transaction on the blockchain reveals the public key that protects other assets on the ledger. This vulnerability pertains to addresses that hide public keys behind a cryptographic hash, namely P2PKH, P2SH, P2WPKH and P2WSH, and occurs when some but not all bitcoin at an address is spent. The spending transaction reveals the public key, defeating the protection normally afforded by the cryptographic hash. The vulnerability persists until all bitcoin locked to the address is spent. Thus, the vulnerability arises from reuse of Bitcoin addresses.
3. **Public Mempool Exposure** vulnerability comes about when the unlocking script exposes the public key of a wallet that is spending bitcoin by revealing the key while broadcasting the transaction to the public Bitcoin mempool. Assuming no reuse or other prior exposure of the public key, the bitcoin remains vulnerable only while it awaits settlement in the public mempool, i.e. until it is permanently recorded on the blockchain. At present, all bitcoin is vulnerable to Public Mempool Exposure or reorganization attacks.
4. **Offchain Exposure** vulnerability results from public key disclosure as an oblique consequence of Bitcoin’s transaction settlement protocol. For example, the owner may have spent digital currency from the public key on another blockchain, such as Bitcoin Cash or Rootstock, or they may have elected to share their public key with a third party, such as a portfolio tracking app. Offchain Exposure vulnerability arises from user practices and business mechanisms built on top of Bitcoin.

Weak Address and Address Reuse vulnerabilities are two types of Onchain Exposure vulnerability [27]. Public Mempool Exposure vulnerability enables only on-spend attacks. By contrast, Onchain and Offchain Exposure vulnerabilities enable both on-spend and at-rest attacks. The latter can be launched using any CRQC, including slow-clock devices. Table I summarizes quantum vulnerabilities of the most prominent existing and proposed script types while Figures 6 and 7 quantify the scale of affected assets.

To understand how Mempool and Onchain Exposure vulnerabilities arise in practice, consider a first time owner who receives their initial bitcoin at an address with the bc1p prefix. These addresses represent P2TR locking scripts that directly record the public key on the blockchain in the receive transaction causing an Onchain Exposure vulnerability that places the user at immediate risk of both on-spend and at-rest attacks from the moment they receive bitcoin. In contrast, a first time owner who receives their initial bitcoin at a previously unused address with the bc1q prefix enjoys potential immunity to at-rest attacks. These addresses represent P2WPKH and P2WSH locking scripts that hide the public key behind a fingerprint, so as long as the user does not share their public key, e.g. by spending some of the bitcoin, there is no public key for an attacker to target with a quantum computer and the user’s asset is immune to at-rest attacks. In particular, their bc1q address may continue to receive bitcoin. However, any transaction that spends some, but not all bitcoin at this address introduces Address Reuse vulnerability. From this point on, all bitcoin still locked to this address is vulnerable to both on-spend and at-rest attacks.

Thus, the presence of Address Reuse vulnerability simplifies, but also darkens, this complex situation. If the nominally safer bc1q address has been used previously, its public key is already present on the blockchain in an older transaction that an attacker can easily identify from the public key’s has value. Thus, transactions recorded on the blockchain act as a “cheat-sheet” that enables an attacker to solve the otherwise completely intractable problem of computing the public key from its hash. Consequently, public key reuse, and indeed any other public key exposure, renders all types of locking scripts equally vulnerable to quantum attacks: the safer scripts, such as P2PKH and P2WPKH, become as vulnerable as the weaker scripts, such as P2PK and P2TR.

Offchain Exposure vulnerabilities arise from user practices and community conventions. Privacy and security recommendations dating back to Satoshi’s original Bitcoin whitepaper [140] dictate that users generate a new address, and hence a new pair of public and private keys, for every transaction. These guidelines are critical in a world with

Script type	Quantum vulnerability	Address prefix	Remarks
P2PK	onchain mempool offchain	n/a	Includes Satoshi era mining rewards
P2PKH	mempool offchain	1...	
P2MS	onchain mempool offchain	n/a	legacy multi-signature script type supplanted by P2SH, P2WSH, and P2TR
P2SH	mempool offchain	3...	
P2WPKH	mempool offchain	bc1q...	currently most popular script type
P2WSH	mempool offchain	bc1q...	
P2TR	onchain mempool offchain	bc1p...	most recently introduced script type
<i>P2MR</i>	<i>mempool</i> <i>offchain</i>	<i>bc1z...</i>	<i>at-rest attack mitigation proposal under discussion</i>

Table I. Quantum vulnerabilities of existing (roman) and proposed (*italics*) Bitcoin script types. Every existing Bitcoin transaction type is vulnerable to on-spend quantum attacks by a potential future fast-clock CRQC, such as one based on superconducting qubits. In addition, bitcoin locked using old P2PK scripts (including ~ 1.7 million bitcoin in Satoshi era mining rewards) and bitcoin locked using modern P2TR scripts (which moved ~ 16.8 million BTC in 2025 and which represented 21.68% of all Bitcoin transactions in 2025) are vulnerable to at-rest attacks by potential future slow-clock CRQCs (per TX outputs in blocks 877259 through 930340). Custom script types, including the unspendable `OP_RETURN` (which can be used to permanently destroy coins or to embed arbitrary data on the blockchain), are not listed.

CRQCs if migration to PQC has not yet occurred, especially in Scenario 2. Indeed, these recommendations, when extended with the rule to only use SegWit addresses (corresponding to P2WPKH and P2WSH scripts) and thus to refrain from using Taproot addresses (corresponding to P2TR scripts), can protect users from at-rest attacks which are the only attacks that a slow-clock CRQC can launch. However, the core recommendation against public key reuse pre-dates concerns about quantum cryptanalysis, because even in a world without quantum computers public key reuse weakens the user’s security (e.g., by exposing them to implementation flaws and side-channel attacks) and privacy (e.g., by facilitating transaction linking, wealth estimation, and deanonymization).

Nevertheless, Bitcoin addresses are often reused for convenience. For example, merchants and exchanges may prefer to publish a single destination address that is stable and recognizable. In fact, public key reuse is extremely pervasive. We quantify the exposure of the largest holdings in Figure 7 where all assets are vulnerable due to public key reuse except those locked by P2PK, P2MS, and P2TR scripts (where public keys are exposed even without address reuse). More detailed and comprehensive data is available in the “Project 11 Risq List” [146] which is a repository of public keys at risk of at-rest attacks. Some of the large holdings in Figure 7 are linked to Binance, Robinhood and Bitfinex, which are also three of the largest cryptocurrency exchanges. In fact, many cryptocurrency services are guilty of address reuse [147]. Reusing public keys helps financial service providers avoid logistical challenges of endless account whitelisting, makes it easy to prove reserves, and helps to reduce fees that they pay in transactions. The key reuse problem also manifests if one uses the same private key on two or more blockchains (e.g., Bitcoin and Rootstock) or if the original blockchain is forked (e.g., Bitcoin and Bitcoin Cash). Spending from an address on one blockchain enables at-rest attacks on the address on all blockchains.

A convenient way of addressing some of the practical challenges of avoiding key reuse is presented by modern Hierarchical Deterministic (HD) wallets [148] which make it easy for users to automatically generate a large number of new keys and addresses in a hierarchical structure. At the top of the hierarchy is a secret phrase called a “seed” which an HD wallet uses to derive so-called extended keys which can in turn be used to derive further extended keys as well as regular keys used for cryptocurrency transactions. There are two types of key derivations: normal and hardened. In both cases, private keys can only be derived from the parent extended private key. However, normal derivation allows public keys to be derived in two ways: from the parent extended private key and from its corresponding extended public key. By contrast, hardened derivation allows public keys to only be derived from the

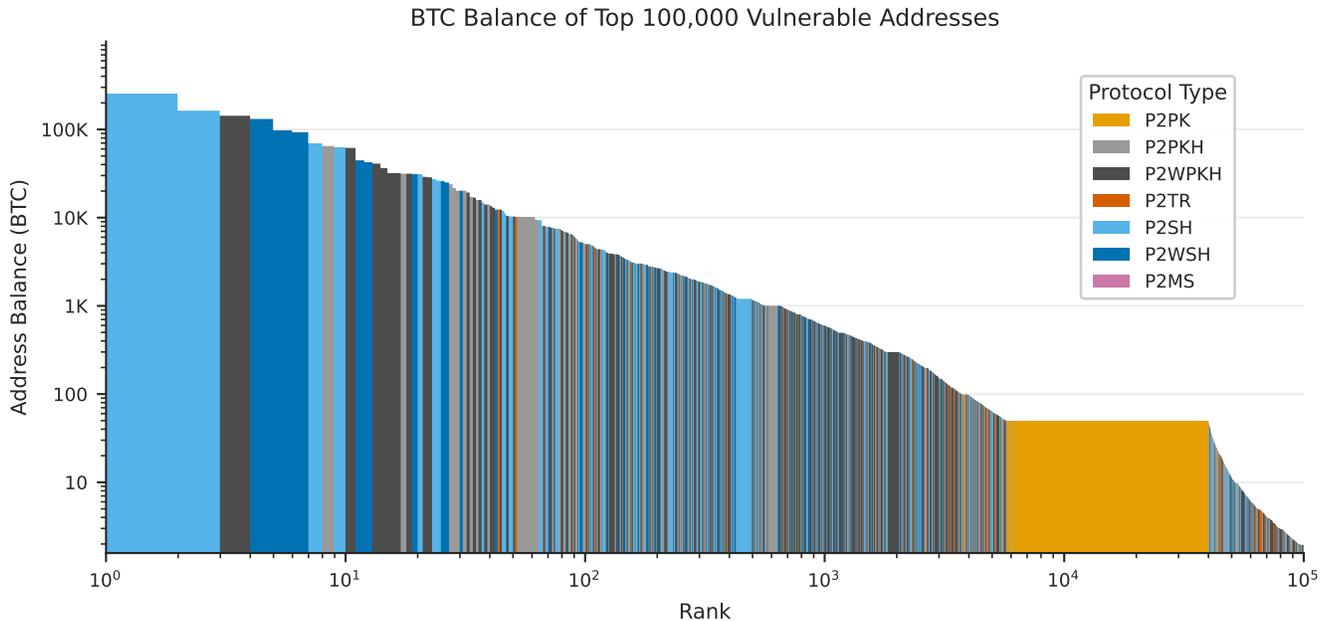


Figure 7. BTC Balance of Top 100,000 Vulnerable addresses. The graph displays the BTC balance of the top 100,000 Bitcoin addresses ranked by value that are vulnerable to at-rest attacks from exposed or re-used keys. In the case of P2PK, P2TR etc. this vulnerability comes from publishing the public key, In the case of other protocols this comes from public key re-use. At around address rank 6000 there is a very large portion of many exposed 50 BTC addresses from the early P2PK mining era because that amount was the mining reward at that point in time. The sum of all BTC in vulnerable addresses is approx. 6.7 million BTC. Plot generated using data from `bigquery-public-data.crypto.bitcoin` [139].

extended private key.

In the absence of CRQCs, it is safe to share extended public keys and use normal derivation to obtain keys for cryptocurrency transactions. This can be used for example to enable a third party service to independently derive a large number of a given user’s public keys without access to their corresponding private keys — a useful feature that enables automatic monitoring of many Bitcoin addresses. Nevertheless, even without quantum computers, normal derivation allows an attacker to recover the extended private key, and hence any derived private key, from the knowledge of a single derived private key and the extended public key. Therefore, in the presence of CRQCs, reuse or disclosure of an extended public key constitutes a pernicious form of Offchain Exposure vulnerability that enables an attacker with a CRQC to compromise multiple private keys at the cost of one. Third party services, such as wallets, portfolio monitoring apps etc., that require sharing of extended public keys or that build centralized repositories of extended public keys potentially aggravate quantum vulnerabilities in the ecosystem.

Indeed, much of the decentralized financial ecosystem built on top of the Bitcoin blockchain was constructed under the assumption that ECDLP-based public keys are safe to share and reuse. This assumption fails in the presence of CRQCs. At the same time, the system’s size and complexity reflect the fact that this assumption often helps to address genuine business needs. However, the possibility of Scenario 2 materializing in near future underscores the importance of greater caution in user and business practices, e.g., by reducing or eliminating reliance on address reuse. These measures will provide a modicum of quantum safety to security-conscious Bitcoin users who are concerned about the possibility of attacks by slow-clock CRQCs and who are willing to eschew the convenience of financial services based on public key reuse. However, such mitigation measures are likely to be seen as insufficient for the broader Bitcoin community. Therefore, the only real long-term remedy to Bitcoin’s quantum vulnerabilities lies in migrating to PQC keys for signing transactions.

D. The Infeasibility of Quantum Attacks on Proof-of-Work

A significant portion of literature addressing quantum vulnerabilities in Bitcoin seems to touch on the topic of using Grover’s algorithm to accelerate mining [25, 149–151] and informal discourse in social and trade media indicates a persistent belief that quantum computers may pose a threat to Bitcoin’s Proof-of-Work consensus mechanism [152,

[153]. However, we maintain that this is not something to worry about in the next several decades.

Grover-based attacks on Bitcoin mining are not practically relevant for two reasons. First, the quadratic quantum speedup from Grover’s algorithm is all but consumed by the overheads of quantum error correction [154]. Second, Grover’s algorithm does not parallelize well [155]. In Bitcoin mining, hardware acceleration and massive parallelization are a much greater advantage than the modest speedup from Grover’s algorithm.

Under the fantastical assumption that a CRQC could compute SHA-256 in a single 1 microsecond error correction cycle, a quantum miner could achieve a hashrate of $0.25 TH/s$ which is more than two orders of magnitude below $110 \pm 3\% TH/s$ of the popular ASIC miner S19 Pro [156]. Under realistic assumptions, rather than fantastical ones, quantum miner’s hashrate plummets by over 10 orders of magnitude [157].

Of course, the long arc of technology is difficult to predict. Quantum computers enable efficient simulation of quantum systems at scale, accelerating innovations in materials science. Among the many fields of technology that stand to benefit from this new capability is the development of quantum hardware itself. Thus, in the long term it is possible that currently unanticipated new quantum architectures may emerge which might allow future generations to build quantum computers with very fast error correction cycles on which Grover speedup becomes relevant in Proof-of-Work mining. But for now, quantum mining remains science fiction more than a concrete threat.

Thus, CRQCs pose no direct threat to Bitcoin mining. However, they do pose an indirect one: if the Bitcoin community does not prepare for the arrival of CRQCs in time and quantum attacks on Bitcoin transactions lead to rapid decline in the value of the cryptocurrency, then the fiat value of mining rewards might change faster than the network is able to adjust mining difficulty, potentially rendering mining unprofitable. This could push some mining pools to shut down a portion of their capacity increasing the average block interval. This would reduce transaction rate which would, in turn, affect user experience and make on-spend attacks easier. These issues arise from Bitcoin’s fixed mining difficulty adjustment schedule which recomputes difficulty targets once every 2016 blocks (about once every two weeks). By contrast, Bitcoin Cash adjusts mining difficulty after every block [158].

IV. QUANTUM VULNERABILITIES OF OTHER ECDLP-BASED CRYPTOGRAPHIC PROTOCOLS

ECDLP-based digital signatures are a basic cryptographic protocol that introduces quantum vulnerabilities into the digital economy. However, many blockchains, including Ethereum, also take advantage of more advanced ECDLP-based cryptographic primitives, such as signature aggregations, key exchange protocols, commitment schemes, and zero-knowledge arguments, in order to enhance privacy and improve scalability. Here, we give a synopsis of these cryptographic protocols and the novel quantum vulnerabilities they introduce. We also discuss a few special elliptic curves and how they influence quantum vulnerabilities.

In certain applications, such as Ethereum’s Proof-of-Stake consensus mechanism, a system may need to aggregate hundreds or thousands of digital signatures into a few dozen. This can be accomplished using ECDLP-based signature aggregation protocol, known as the Boneh-Lynn-Shacham (BLS) scheme [159]. BLS signatures depend crucially on the ability to efficiently compute a certain non-degenerate bilinear map, called pairing, and therefore require specialized elliptic curves, called pairing-friendly curves, for which pairing can be efficiently computed. This is impossible on the secp256k1 curve, used for most digital signatures in Bitcoin and Ethereum, so protocols that rely on pairings utilize other curves, such as `alt_bn128` [160] from the Barreto-Naehrig curve family [161] and BLS12-381 [162] from the Barreto-Lynn-Scott family [163].

Key exchange protocols enable two parties in direct communication over a public channel to establish a shared secret that passive eavesdroppers cannot recover from the exchanged messages. Diffie-Hellman (DH) key exchange [164] and its variant based on elliptic curves called Elliptic Curve Diffie-Hellman (ECDH) key exchange are the most well-known examples. Both DH and ECDH are based on the discrete logarithm problem and allow passive eavesdroppers with a CRQC to recover the shared secret, breaking the privacy normally afforded by the protocol. ECDH is used in certain privacy-preserving blockchains, such as Zcash [165], Monero [166], and Litecoin’s Mumblewimble [167].

Commitment schemes are cryptographic protocols that allow one to publish a value that commits them to a choice without revealing that choice. Such schemes have two properties: the choice cannot be altered (binding) and remains secret (hiding). Homomorphic commitments have the additional property that the sum of individual commitments to a set of values equals the commitment to the overall sum. In the context of cryptocurrency transactions, such commitment schemes can be used to hide transaction amounts while enabling the network to verify that inputs and outputs of each transaction balance correctly. Popular examples of homomorphic commitments based on ECDLP are ElGamal [168] and Pedersen [169] commitment schemes. Quantum attacks on the former destroy the hiding property and on the latter destroy the binding property [170, 171]. Consequently, if such a commitment scheme is used to hide transaction values on a privacy-preserving blockchain, a quantum attacker can break the privacy protection in the former case or create new coins in the latter case. An example of a more sophisticated commitment scheme is the Kate-Zaverucha-Goldberg (KZG) protocol [172] which allows a party to commit to a polynomial and subsequently

to provide efficient openings that reveal evaluations of the polynomial. KZG scheme is currently used for example in Ethereum’s DAS mechanism [173].

Related cryptographic protocols are the Zero-Knowledge (ZK) [41] proofs and arguments, including those we use in this paper as the means of verifying quantum resource estimates without disclosing sensitive attack details. These protocols allow one party to convince another that a statement is true without revealing the facts that make it true. In addition to this zero-knowledge property, a secure ZK protocol is supposed to be complete (proofs and arguments of true statements pass verification) and sound (proofs and arguments of false statements are almost guaranteed to fail verification). The distinction between proofs and arguments concerns the computational power of the proving party: in a proof the prover is assumed computationally unbounded while in an argument the prover’s computational resources are limited [174]. If the statement to be proven takes the form “prover knows secret X”, then the protocol is called a proof or argument of knowledge. ZK protocols are typically interactive and involve one party challenging the other to compute values that depend on, but do not reveal, the secret. However, the Fiat-Shamir heuristic [91] allows one to turn an interactive protocol into a non-interactive one. This transformation is a core component of efficient ZK protocols known as Zero-Knowledge Succinct Non-interactive Arguments of Knowledge (zkSNARKs) [175]. Non-interactive protocols that possess the succinctness, soundness and completeness property, but not necessarily the zero-knowledge property, are called SNARKs. In blockchain technologies, zkSNARKs can be used to enhance privacy while SNARKs (zero-knowledge or otherwise) can be used to achieve scalability and improve efficiency.

Many commitment schemes and ZK protocols rely on the assumption that nobody knows certain secrets related to the protocol’s fixed public parameters. For some schemes, such as the KZG polynomial commitments, the secrets, sometimes referred to in this context as “toxic waste”, are known temporarily before being erased during the trusted setup ceremony [69] in the course of which the protocol parameters are first established. For other schemes the parameters are produced by methods satisfying the Nothing Up My Sleeve (NUMS) principle [176–178] to ensure that nobody learns them in the process of choosing protocol parameters. However, for certain ECDLP-based protocols the secrets can be derived from the fixed public parameters using a CRQC. Moreover, since protocol parameters are fixed, this needs to be done only once and the resulting secrets can be reused to launch multiple attacks later using classical computers. Example protocols vulnerable to such on-setup attacks include Pedersen commitments used in Mumblewimble, KZG commitments used in Ethereum’s Data Availability Sampling mechanism, and BulletProofs [179] used in Monero and Mumblewimble.

Consequences of a quantum attack on an ECDLP-based SNARK generally depend on which cryptographic property of the underlying commitment scheme is quantum vulnerable. For instance, a quantum break of the commitment’s binding property collapses the soundness property of the SNARK, allowing forged validity proofs for unauthorized transactions to pass verification. Moreover, in privacy-preserving settings such attacks can remain undetected until systemic effects materialize. For example, a CRQC enables on-setup attacks on the Tornado Cash protocol in which attackers withdraw more funds than they deposited. Due to the privacy properties of the protocol, they can drain an anonymity pool without anybody noticing until the pool’s balance falls to zero.

The succinctness of SNARKs can be used to build Layer 2 (L2) protocols [180] on top of a Layer 1 (L1) blockchain, such as Ethereum, to increase transaction rate and reduce fees. L2 protocols that use SNARKs are known as zk-rollups. They process a large number of transactions offchain and bundle them together into batches that are subsequently committed to the underlying L1 blockchain using SNARKs. Quantum vulnerabilities in these protocols potentially enable attackers to convince the onchain verifier that a batch of L2 transactions is valid even when it includes invalid transactions that, for example, steal assets or inflate monetary supply. In the long term, these vulnerabilities may be remediated by switching to hash-based or lattice-based protocols which are regarded as post-quantum secure. Indeed, some zk-rollups, such as Starknet, use hash-based protocols which are believed to be resistant to quantum cryptanalysis.

Quantum-vulnerable SNARKs typically employ ECDLP-based commitment schemes which, like BLS signatures, require pairing-friendly elliptic curves. Our quantum resource estimates discussed earlier apply directly to solving ECDLP on curves such as secp256k1 which is not pairing-friendly. The estimates are sensitive to curve parameters and do not apply unchanged to all elliptic curves used in blockchain cryptography. We have not conducted rigorous resource estimates for ECDLP on other elliptic curves, such as those used in SNARKs. However, we expect the cost of solving ECDLP on many curves over finite fields with 256-bit modulus and 256-bit group order to be on the same order of magnitude. This includes the cost of ECDLP on the Pasta curves [181] employed in zkSNARKs behind the newest type of shielded (private) transactions in Zcash. We expect the cost to be only moderately higher for ECDLP on other 256-bit curves such as alt_bn128 [160] whose efficient implementation is provided as precompiled smart contracts [182] on Ethereum and used for example in zkSNARKs behind its cross-chain integrations and zk-rollups. However, ECDLP on curves over larger finite fields may require additional quantum memory. In particular, the BLS12-381 curve [162, 163] used in older types of shielded transactions in Zcash and in Ethereum’s Proof-of-Stake consensus mechanism, requires 381 bits to store point coordinates. This leads to a 50% increase in the size of some quantum registers and implies that solving ECDLP on BLS12-381 may require a somewhat larger CRQC than the

above 256-bit curves. Nevertheless, because of the efficient scaling of Shor’s algorithm [46], we still expect ECDLP on this curve to be accessible to the first CRQCs.

The amount of time and the number of qubits that an early CRQC will need to solve ECDLP on a specific elliptic curve increases with the modulus and the order of the elliptic group. Asymptotically, this increase is modest (merely polylogarithmic) due to the efficiency of Shor’s algorithm. However, the question of whether increasing the modulus and the group order are appropriate stopgap measures to protect ECDLP-based cryptosystems against attacks by early CRQCs is more complex. For example, solving ECDLP on an elliptic curve with 1024-bit modulus and group order may require a CRQC with almost five thousand logical qubits. Moreover, quadrupling the number of bits in the modulus and group order will increase the runtime of Shor’s algorithm by roughly a factor of 64, although (as mentioned before) the runtime can be reduced if the number of qubits is increased. Similarly, an attack on the strongest ECDLP-based cryptography in the TLS protocol, which is used to encrypt and authenticate HTTPS traffic and which currently supports a 521-bit elliptic curve [16], may require a slightly larger CRQC than an attack on the BLS12-381 curve and the other curves discussed here.

However, the question whether a switch to a large, e.g. 1024-bit, modulus can provide blockchains with a temporary protection is complex and depends on a detailed understanding of scaling barriers in leading quantum computing platforms within relevant machine sizes. Given broad progress across multiple hardware architectures, the safe assumption is that there may be little time between the breaking of 256-bit ECDLP and the breaking of 1024-bit ECDLP. Therefore, large elliptic curves may provide no real protection against at-rest attacks. At the same time, such curves would, at least initially, increase the time needed to derive private keys, so they may provide a short reprieve from on-spend attacks. In any case, we expect the reprieve to shrink as CRQCs become larger and are able to trade quantum memory for time. Thus, the security benefits of large elliptic curves are, at best, partial and temporary and, at worst, nearly non-existent.

V. QUANTUM VULNERABILITIES OF THE ETHEREUM BLOCKCHAIN

Ether (ETH), is the native digital asset stored and traded on the Ethereum blockchain. It is the world’s second largest cryptocurrency by market capitalization (approximately 400 billion USD as of February 2026), second only to Bitcoin (approximately 1.9 trillion USD as of February 2026). However, relying solely on cryptocurrency market capitalization obscures the full economic weight of the Ethereum network. When measured by the Total Value Secured (TVS)[183], which sums the native ether, over 150 billion USD in fiat-pegged stablecoins, tokenized Real World Assets (RWAs), and the fast-growing Layer 2 ecosystem, Ethereum’s aggregate economy amounts to over 600 billion USD. In contrast, Bitcoin’s ecosystem (including Layer 2s like Lightning Network [184] and tokens) adds a relatively modest ~15 billion USD to its base value, resulting in a TVS still largely dominated by the 1.9 trillion USD market cap of the cryptocurrency itself.

While Bitcoin is a world-wide distributed ledger recording transactions in a single native digital asset, Ethereum is effectively a world-wide distributed computer that employs blockchain technology to store value and information [185, 186]. It does not merely record balances, but executes code which implements a wide variety of financial, semi-financial and non-financial applications, such as derivatives, hedging contracts, savings accounts, wills, reputation systems, peer-to-peer gambling applications, prediction markets, non-fungible tokens (NFTs) and online voting systems [185]. By creating trust in a trustless environment through code and cryptographic protocols rather than through centralized institutional intermediaries, the system enables new forms of financial arrangements, such as decentralized exchanges and peer-to-peer lending, known as Decentralized Finance (DeFi). Ethereum smart contracts [186, 187], which enable realization of these diverse applications, are sufficiently powerful to facilitate implementation of the bylaws of an entire business organization in the form of a stateful long-term smart contract sometimes referred to as a Decentralized Autonomous Organization (DAO) [185]. The combination of smart contracts and blockchain technology creates a smart legal system for a digital economy in which smart contracts describe and enforce the rules and the blockchain acts as a decentralized notarial service providing a shared source of truth about the economy’s current state. Ethereum’s grand ambition is to become “a decentralised secure social operating system” that enables business contracts to “algorithmically specify and autonomously enforce rules of interaction” [186].

Ethereum’s quantum security posture is very different from Bitcoin’s. For example, Ethereum produces new blocks in deterministic 12-second slots with most transactions processed in less than a minute (see Figure B.3 in [188]). Therefore, our estimates indicate that early fast-clock CRQCs are unlikely to be able to launch on-spend attacks against Ethereum. Moreover, the Ethereum ecosystem has sophisticated private mempools, including TEE-based BuilderNet, that allow users to submit transactions directly to builders which provides a mitigation for potential on-spend attacks.

The main quantum threat against Ethereum lies in a variety of at-rest attacks. The Ethereum blockchain has three prominent features differentiating it from Bitcoin that give rise to five types of quantum vulnerabilities. First is

Ethereum Component	Vulnerability	Vulnerable Cryptographic Primitives	Assets at Risk	Second Order Effects
Account model	Account Vulnerability	ECDSA	User funds: 20.5M ETH	Compromise of Hot wallets, store accounts etc.
Smart Contracts	Admin Vulnerability	ECDSA	Contract TVS: 2.5M ETH, 200B USD in stable-coins/RWAs	Compromise of Oracles, RWAs, Bridges, Guardians etc.
	Code Vulnerability	ECDSA, alt_bn128, KZG, BLS12-381 [182]	L2/Protocol TVS: 15M ETH	Compromise of L2s with need for admin intervention
Validators	Consensus Vulnerability	BLS Signatures	Consensus stake: 37M ETH	Compromise trust in the blockchain itself
	Data Availability Vulnerability	KZG Commitments	L2/Protocol TVS: 15M ETH	Compromise trust in the blockchain itself

Table II. Taxonomy of Ethereum Components and Vulnerabilities. This table summarizes the vulnerabilities, cryptographic root causes, assets at risk and second order effects from the subsequent sections. All of these vulnerabilities are at-rest attacks.

the account model: unlike Bitcoin, where currency units exist as UTXOs, Ethereum maintains global state through persistent accounts. While users can rotate addresses, certain ecosystem incentives (maintaining DeFi positions, governance history and stable identifiers) can incentivise the use of longer lived accounts [186]. Once an account initiates a transaction, its public key remains exposed indefinitely. Second is the native support for smart contracts which are autonomous computer programs executed by the Ethereum Virtual Machine (EVM) that can hold assets and enforce complex rules without intermediaries [186]. Third is the Proof-of-Stake (PoS) consensus mechanism in which Sybil resistance based on energy-intensive mining is replaced with a system where “validators” secure the network by staking their own capital as collateral [189–191].

These three architectural choices introduce five distinct attack vectors: the Account Model creates the Account Vulnerability (due to the use of quantum-vulnerable ECDSA [192]); Smart Contracts introduce Admin Vulnerability [193] and Code Vulnerability; and the specific cryptography required for validators gives rise to the Consensus Vulnerability (due to the use of BLS signatures [159]) and Data Availability Vulnerability (due to the use of KZG commitments [172]). In the following sections, we will define each type of vulnerability, quantify risk (in native ETH and other assets), estimate the cost of quantum attacks and review some mitigation measures. We summarize these vulnerabilities in Table II.

A. Account Vulnerability

Bitcoin prioritizes privacy by tracking UTXOs rather than persistent user identities. Unlike Bitcoin where address rotation can avoid at-rest vulnerabilities, Ethereum’s account model is structurally prone to at-rest attacks. There are two types of accounts: Externally Owned Accounts (EOA), controlled by a private key and possessing agency to initiate actions on the blockchain, and contract accounts, governed by code and able to react to transactions sent to their address [186]. Unlike Bitcoin, Ethereum users typically maintain the same EOA for an extended period of time to accumulate assets and reputation (identity). Because an EOA’s address is derived directly from its public key, the keys cannot be rotated; securing an account with a new key pair requires abandoning the account entirely [186, 194]. Crucially, while an address initially masks the public key, the moment a user sends their first transaction, the digital signature reveals the full public key to the network [192]. This results in Account Vulnerability: a systemic, unavoidable exposure that cannot be mitigated by user behavior, short of a protocol-wide transition to PQC. We estimate that a quantum attacker using a fast-clock CRQC could crack the 1,000 highest-net-worth Ethereum accounts (at the time of writing, holding approximately 20.5 million ETH, see Figure 8) in less than nine days.

Account Vulnerability has other potential consequences. For example, it allows attackers to forge votes in decentralized governance of Decentralized Autonomous Organizations (DAOs). It also affects Exchange Hot Wallets: automated, internet-connected accounts used by centralized exchanges to process high throughput withdrawals. Even though ETH balance in these wallets may be low at a given point in time, the overall danger is very high: a breach here does not just represent a loss of funds for the exchange, but could lead to the draining of custodial liquidity that could induce panic across a broader market.

The Ethereum community has deployed several proposals that mitigate Account Vulnerability, most notable of which is Account Abstraction (AA) via ERC-4337 [194]. Introduced in 2023, AA enables users to interact with

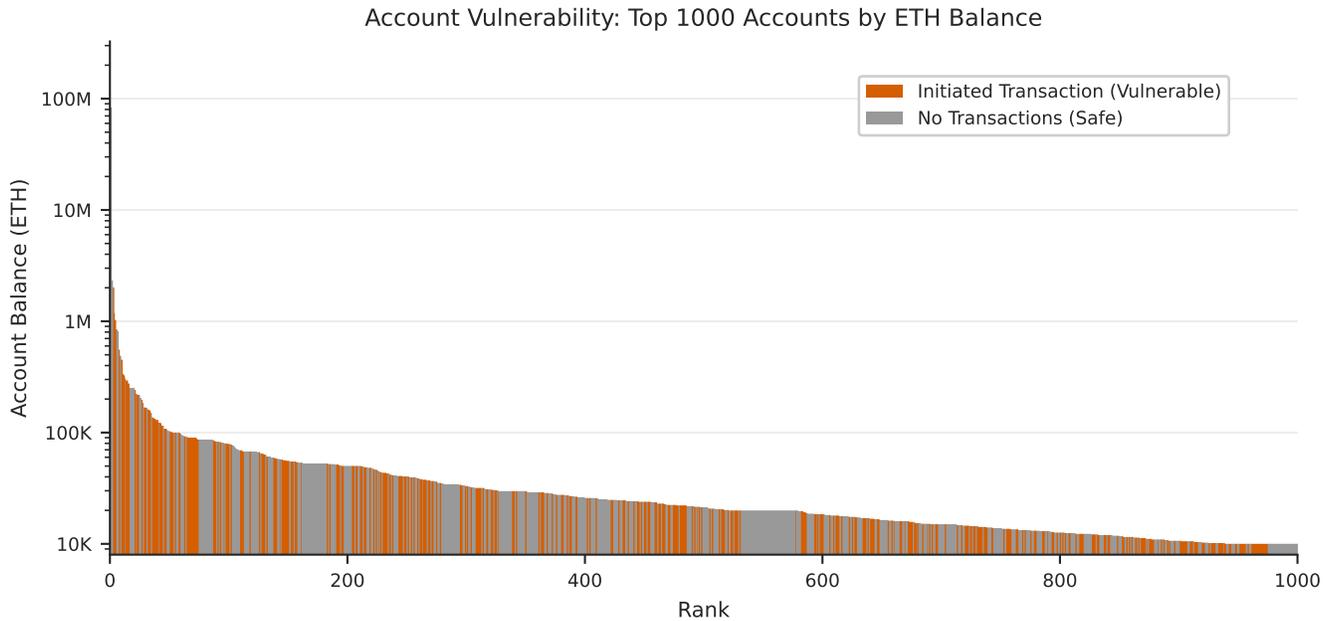


Figure 8. Account Vulnerability of Top 1000 accounts by ETH balance. The graph displays the ETH balance of the top 1000 Ethereum accounts ranked by value. Orange indicates accounts that have initiated a transaction and are therefore vulnerable to at-rest attacks, grey accounts have not. The sum of all ETH in vulnerable accounts in this list is ~ 20.5 million ETH. Plot generated using data from `bigquery-public-data.crypto_etherium` [195].

Ethereum via smart wallets that offer (among other features) customizable authentication logic. This decouples user identity (reputation) from a single static key, reducing the quantum attack surface through more frequent key rotation. Furthermore, AA helps mitigate the future challenge of migrating the “long tail” of legacy accounts to PQC without requiring a protocol-level hard fork. More recently, EIP-7702 [196] expanded these capabilities by allowing EOAs to temporarily function as smart contracts. However these enhancements mitigate the symptoms and not the root cause: the only gateways for external agency into Ethereum are the vulnerable EOAs. Thus, while these upgrades improve flexibility, they fall short of providing a fully quantum-secure interface.

B. Admin Vulnerability

An important distinction between Bitcoin and Ethereum is the latter’s extensive support for smart contracts [185]. While Bitcoin scripts are fundamentally stateless, Ethereum smart contracts are fully fledged distributed applications with persistent state. They can initiate transactions, access blockchain data, and even autonomously deploy other contracts [186]. To facilitate the ability for different smart contracts and decentralized applications to interact with one another, the Ethereum community established standardized application programming interfaces known as Ethereum Request for Comments (ERCs). These include: ERC-20 [197] for fungible tokens, ERC-721 [198] for Non-Fungible Tokens (NFTs), and ERC-3643 [199] for compliance-aware Real World Assets (RWAs). These standards have helped fuel a massive “tokenized” economy that resides on top of Ethereum but is distinct from the native ether (ETH) asset (see Figure 10). Furthermore, the rapid adoption of digital representations of bonds, stocks, commodities, etc., as tokens [200] means that the value exposed to Ethereum’s security model extends far beyond the native currency.

Smart contracts are a source of two quantum vulnerabilities. The first one, which we term the Admin Vulnerability, arises from Account Vulnerability of accounts for which a smart contract reserves enhanced privileges or administrative control, such as the ability to pause execution, upgrade code or extract funds [193]. Because these administrative keys are rarely rotated and often publicly utilized for governance voting or contract upgrades, they represent a high-value at-rest vulnerability. At the time of writing, among the top 500 contract accounts by ETH balance, at least 70 accounts, with total holdings of about 2.5 million ETH, see Figure 9, are at risk of account takeover by a quantum attacker due to Admin Vulnerability. We expect that a private key derivation attack on these 70 Ethereum accounts using a fast-clock CRQC will take less than 15 hours.

A strategic adversary could prioritize “high-leverage” admin accounts that exhibit Admin Vulnerability. These

Admin Vulnerability: Top 500 Contracts by ETH Balance

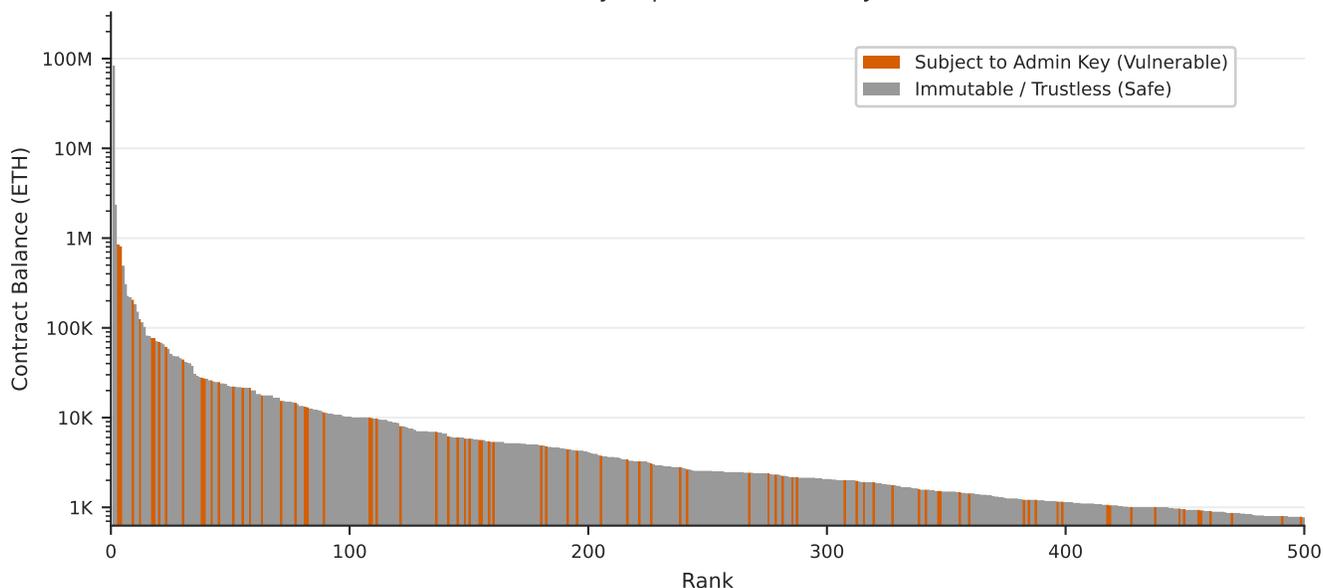


Figure 9. Admin Vulnerability of Top 500 smart contracts by ETH balance. Contracts were classified as “Subject to Admin Key” if their event logs contain signatures corresponding to AdminChanged, Upgraded (ERC-1967 proxy standards), or OwnershipTransferred (OpenZeppelin Ownable standard). The sum of all ETH in these vulnerable contracts is ~ 2.5 million. Plot generated using data from `bigquery-public-data.crypto_ethereum` [195].

accounts often hold negligible amounts of ETH to minimize exposure, but confer various administrative privileges. A successful quantum attack on an admin account could grant the attacker control over the following (non-exhaustive) list of functions and systems:

1. Real World Asset (RWA) token issuance, see [Figure 10](#): RWA administrative keys used to authorize the minting and burning of tokens that represent offchain collateral; compromising them allows an attacker to mint fraudulent tokens to collapse the peg between a digital token and the real-world asset backing it.
2. Bridges: Bridge administrators manage the multi-signature schemes that lock assets on one blockchain to mint “wrapped” equivalents on another; compromising them allows the attacker to drain the entire liquidity pool backing these cross-chain transfers (This risk supersedes the Code Risk we analyze later, see [Figure 11](#). For L2 scaling solutions, admin failure is more catastrophic than an ordinary protocol failure).
3. Oracle Nodes: These accounts broadcast authoritative data feeds (e.g., ETH/USD price) that smart contracts rely on; by hijacking these keys, an attacker could broadcast false price data to trigger catastrophic price action potentially leading to automated liquidations across decentralized lending protocols.
4. Guardians: These privileged accounts act as the emergency safety mechanisms for DeFi applications, holding the power to instantaneously pause execution, freeze user funds, or in some cases, bypass security time-locks to inject code updates.

Quantifying the aggregate financial exposure from Admin Vulnerability as a whole is difficult due to the complex web of dependencies and composability in the Ethereum ecosystem. While the first order 2.5 million ETH (ETH from admin-vulnerable contracts, see [Figure 9](#)) and ~ 200 billion USD in stablecoins and tokenized RWAs (assets tied up in admin-vulnerable contracts, see [Figure 10](#)) is large, the second order risk from other “high-leverage low ETH accounts” (Bridges, Oracles and Guardians etc.) is arguably much larger due to their compromise potentially leading to mass liquidations across lending markets, peg collapse of stablecoins, freezing of cross-chain liquidity etc. This makes the true value at risk very hard to calculate using standard asset-balance models (especially with growth projections of RWAs to 16.1 trillion USD in 2030 [51]), as the effective exposure almost certainly encompasses substantial fractions of the entire ecosystem’s TVS.

Admin Vulnerability: Tokenized Assets on Ethereum

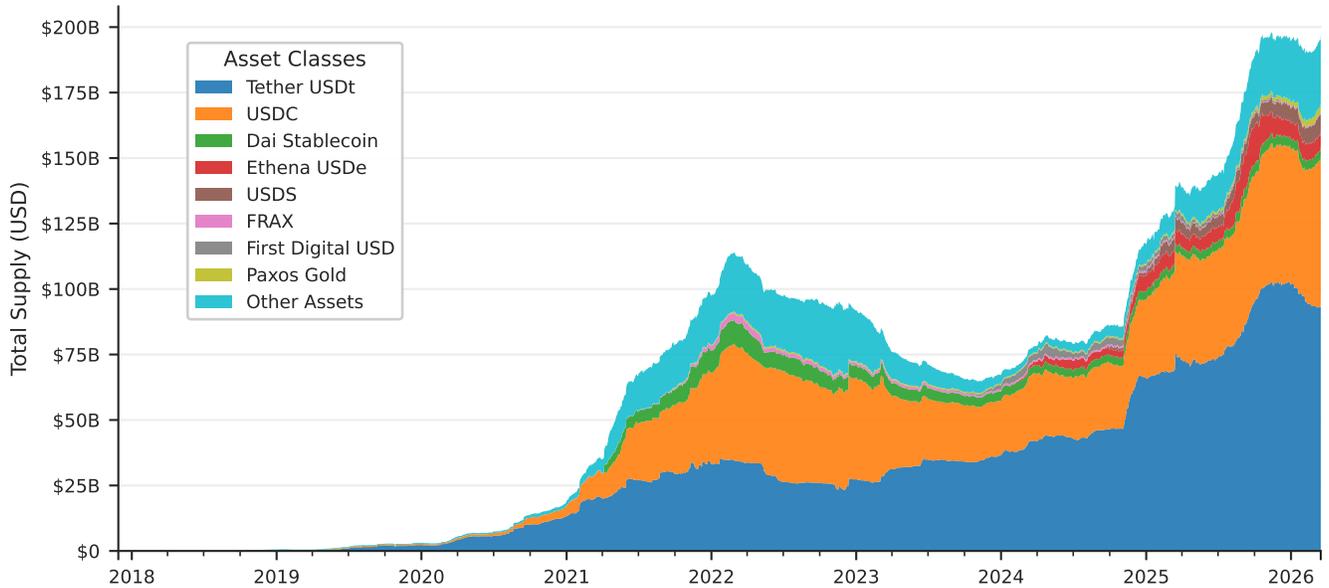


Figure 10. Admin Vulnerability exposure across distributed Real World Assets (RWAs). The chart details the market capitalization of major distributed assets (onchain tokens) on Ethereum, explicitly excluding the underlying offchain represented assets (e.g., fiat reserves, physical properties, repurchase agreements). These distributed assets rely on quantum-vulnerable ECDSA keys for administrative control of their smart contracts. Within this ecosystem, distributed fiat-backed stablecoins and tokenized treasuries present a systemic risk; a compromise of issuer keys would allow arbitrary minting, effectively depegging the distributed asset and destabilizing the broader DeFi economy, even though the offchain represented reserves remain untouched. Private credit protocols carry institutional risk, where attackers could manipulate onchain loan books to redirect repayment flows or bypass credit underwriting standards. Finally, commodities, real estate, and equities embody legal title risk, where an admin takeover would allow the seizure of digital ownership rights, severing the legal link between the distributed token holder and the underlying represented asset. Ethereum currently captures the vast majority of RWA market share (and is therefore the focus of this chart), but it is not the sole network hosting these assets. As this multi-chain landscape evolves rapidly, other blockchain networks must proactively assess their own quantum attack vulnerabilities to protect their respective distributed ecosystems. The landscape of RWAs is evolving rapidly, and so these estimates should be viewed as a point-in-time assessment of the risk profile for various asset classes that will likely become outdated very quickly. Data gathered from: [201].

C. Code Vulnerability

The second smart contract vulnerability concerns the code running on EVM. In order to prevent abuse and runaway algorithms from exhausting system resources, the EVM meters the runtime cost of executing smart contracts [186]. The account sending a transaction is charged the so-called “gas fee” for the cost of executing the smart contracts it triggers. This provides economic incentives for computational efficiency and discourages implementation of complex algorithms, such as cryptographic primitives, directly in EVM byte code (in line with the famous dictum “do not roll your own crypto”). Instead, basic cryptographic primitives, such as hash functions and digital signature verification, are provided as standard precompiled contracts [182, 186] (also known as “precompiles”) that a user’s contract can interact with. However, at present, none of the precompiled contracts implement any of the modern PQC primitives. This situation causes the second quantum vulnerability in smart contracts (also an at-rest vulnerability), which we call Code Vulnerability and which lies in the fact that no precompiles for post-quantum cryptographic protocols, such as zero-knowledge (ZK) arguments, are currently available on Ethereum. Inevitably, the primitives currently utilized by smart contracts carry their own quantum vulnerabilities.

The severity of the Code Vulnerabilities and the challenges of addressing them are magnified by the extensive ecosystem built on top of Ethereum. In the process of becoming the world’s second largest blockchain by market capitalization, Ethereum ran into scalability limits of a single blockchain. Further growth was enabled by development of effective L2 networks, such as state channels and rollups, which process transactions quickly and cheaply offchain and periodically commit transaction batches to an L1 blockchain, such as Ethereum. Assets are moved between the L1 and L2 networks using bridges which rely on event logs, smart contracts, and EOAs to interact with the L1 chain. Therefore, bridges and L2 networks inherit all the vulnerabilities of smart contracts and the Ethereum account

model [202]. In addition, some L2 solutions use quantum-vulnerable zkSNARKs [175, 203]. The Ethereum community is discussing various approaches for introducing PQC at different layers of the Ethereum stack. In particular, EIP-7932 [204] is a proposal for the introduction of precompiles implementing post-quantum signature schemes.

We estimate the Code Vulnerability risk as at least 15 million ETH by summing the TVS of a few major quantum-vulnerable L2 protocols and cross-chain bridges, such as Arbitrum, Base, and Optimism obtained from the L2BEAT dashboard [205], see Figure 11. However, this risk is partially mitigated by the fact that many L2 solutions operate in a “Stage 0” or “Stage 1” maturity phase, retaining centralized administrative safeguards. These security councils typically possess the authority to pause the bridge or override proof verification in the event of an anomaly. Paradoxically, this centralization (while an Admin Vulnerability in its own right) provides a temporary defense against Code Vulnerability; if an attacker were to exploit a cryptographic flaw (e.g., forging a validity proof), human administrators could theoretically intervene to freeze the protocol before the theft is finalized. The eventual upgrade of the Ethereum base layer to quantum-secure consensus does not retroactively fix this Code Vulnerability embedded in existing smart contracts (the EVM cannot automatically recompile deployed smart contracts). To completely remove the Code Vulnerability, the onus falls upon the protocol stakeholders and governance councils to coordinate the migration of their specific bridge contracts and administrative multi-signature schemes. Without this distributed effort, a quantum adversary could derive the private keys for these distinct protocols from public chain data, effectively bypassing the security of the base layer to drain these specific protocols.

D. Consensus Vulnerability

The last major difference between Bitcoin and Ethereum concerns the consensus mechanism. In 2022, seven years after its launch as a Proof-of-Work blockchain, Ethereum upgraded to the energy-efficient Proof-of-Stake consensus

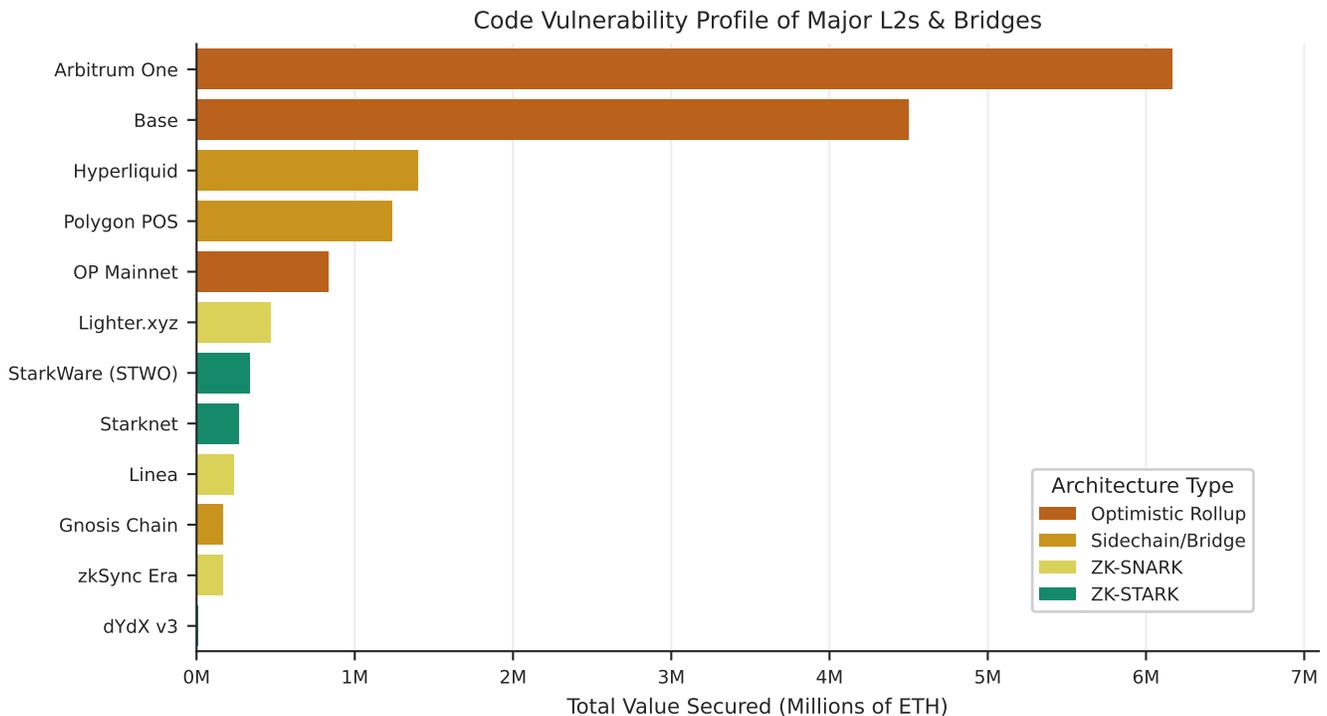


Figure 11. Breakdown of Total Value Secured (TVS) across major scaling protocols, categorized by their underlying security model. Optimistic Rollups (red), Bridges (orange), and zkSNARKs (yellow) currently rely on digital signatures or elliptic curve pairings that are vulnerable to quantum key derivation. Only protocols utilizing zkSTARKs (green) employ hash-based proofs that are currently believed to be resistant to quantum attacks. Across the three examples in each of the four protocol types we chose, the sum of all TVSs is ~15 million ETH. Note: the Layer 2 landscape evolves very rapidly; TVS figures fluctuate daily, and protocols frequently upgrade and change their proving architectures (e.g., migrating from pairing-based to hash-based proofs or hybrid systems that can drastically alter the quantum risk profile). These estimates should be viewed as a point-in-time assessment of the ecosystem’s risk profile that will likely become outdated (or incorrect) very quickly.

mechanism by merging its execution layer with the new consensus layer, called the Beacon Chain, and effectively ceasing all mining activity [189]. In the new protocol, consensus is driven by a one half majority vote among “validators”, with finality established by a two thirds majority [190, 206]. Any Ethereum node can become a validator by staking 32 ETH as collateral. The stake enables the network to incentivize the validator’s honest behavior by a combination of financial penalties and rewards [190].

Penalties suffered by a malicious or misbehaving validator include leaking for inactivity and slashing for provably malicious behavior. Rewards for an actively participating honest validator consist of transaction fees from proposing blocks and issuance from attesting. For most validators, part of the transaction fees come from block builders [207] that extract Maximal Extractable Value (MEV - analogous to the same concept in Bitcoin where it stands for Miner Extractable Value), from their ability to reorder, insert and remove transactions as they build new blocks [143, 144]. In particular, validators’ control over which transactions are executed and when creates opportunities to front-run and sandwich transactions which allows them to benefit financially at the expense of other users of the network [143, 144].

Efficient processing of thousands of digital signatures per second from validators signing attestations presents a cryptographic challenge due to the large volume of cryptographic information involved. Ethereum substantially reduces verification overhead by compressing tens of thousands of digital signatures into a few dozen using the quantum-vulnerable BLS signature aggregation protocol [159] on the BLS12-381 curve. As discussed earlier, Shor’s algorithm on BLS12-381 uses a greater amount of quantum memory than that required by our resource estimates for curves like secp256k1. However, we estimate that the additional resource cost is modest. Consequently, the consensus layer of Ethereum should be considered at-rest vulnerable to the same first-generation CRQCs as the execution-layer account model and smart contracts.

The severity of this Consensus Vulnerability depends on the fraction of validators compromised:

1. If less than a third of validators are compromised, an attacker can force them to “equivocate”, i.e., sign conflicting blocks for the same slot causing the system to slash their stake and eject them from the network [190].
2. If more than a third of validators are compromised, an attacker can prevent finalisation which requires two thirds supermajority. The attacker-controlled stake would dilute via the activity leak until enough of the honest stake regains the supermajority [190]. This constitutes a denial of service attack on finality.
3. If more than half of validators are compromised, an attacker can control the fork choice rule to prevent chain growth or execute deep blockchain reorganizations.
4. If more than two thirds of validators are compromised, the attacker can finalize inconsistent chains.

As of February 2026, about 37 million ETH is staked and exposed to potential slashing and ejection due to quantum attacks exploiting the Consensus Vulnerability, see Figure 12. The second-order risks of Consensus Vulnerability are also very large. First, the deep integration of Liquid Staking Tokens (LSTs) means that large slashing events could trigger a peg collapse for derivative assets widely used as collateral, precipitating cascading liquidations across the DeFi ecosystem analogous to a traditional bank run. Second, a breach of finality introduces large desynchronization risks for cross-chain infrastructure, where an attacker could exploit deep chain reorganizations to double-spend assets across bridges and Layer 2 networks that relied on the immutability of the base layer. Lastly, absolute control over the blockchain and all recent history would likely remove the network’s credible neutrality and possibly the institutional trust required for Ethereum to function as a global settlement layer.

Recovery from a supermajority attack, where more than two thirds of validators are compromised, is impossible within the protocol rules and would necessitate a “social consensus” intervention. The community would be forced to coordinate some emergency measures (possibly a hard fork), manually identifying the last valid block and migrating to a new protocol state; a catastrophic event for trust in the system. Consequences of strategically straightforward quantum attacks on Ethereum consensus layer are primarily destructive in nature: loss of stake inflicted on compromised validators’ accounts, halt of transaction processing. Strategically more sophisticated attacks are possible and potentially involve manipulation of the selection and ordering of transactions in newly proposed blocks of the same character as the one used by validators to obtain the MEV revenue stream [143]. Quantum attackers with a fast-clock CRQCs could also attempt to profit from transaction manipulation by attacking the mempool directly rather than by compromising validators [143].

Attacks on consensus are made easier by the exposure of validators’ public keys in deposit transactions and the validator registry [186, 189]. Another weakness is the absence of expedient mechanisms for key rotation: a validator must go through a lengthy process of withdrawal and re-staking [190]. At the same time, the very large number (about one million) of validators together with the stake-weighted random selection process makes attacks on consensus hard and expensive [190]. Indeed, there is safety in numbers: even assuming that an early fast-clock CRQC can solve ECDLP on BLS12-381 in the same amount of time as on secp256k1, a quantum attacker with 20 such machines would need more than nine months to derive private keys of a two thirds supermajority of validators. It should be noted

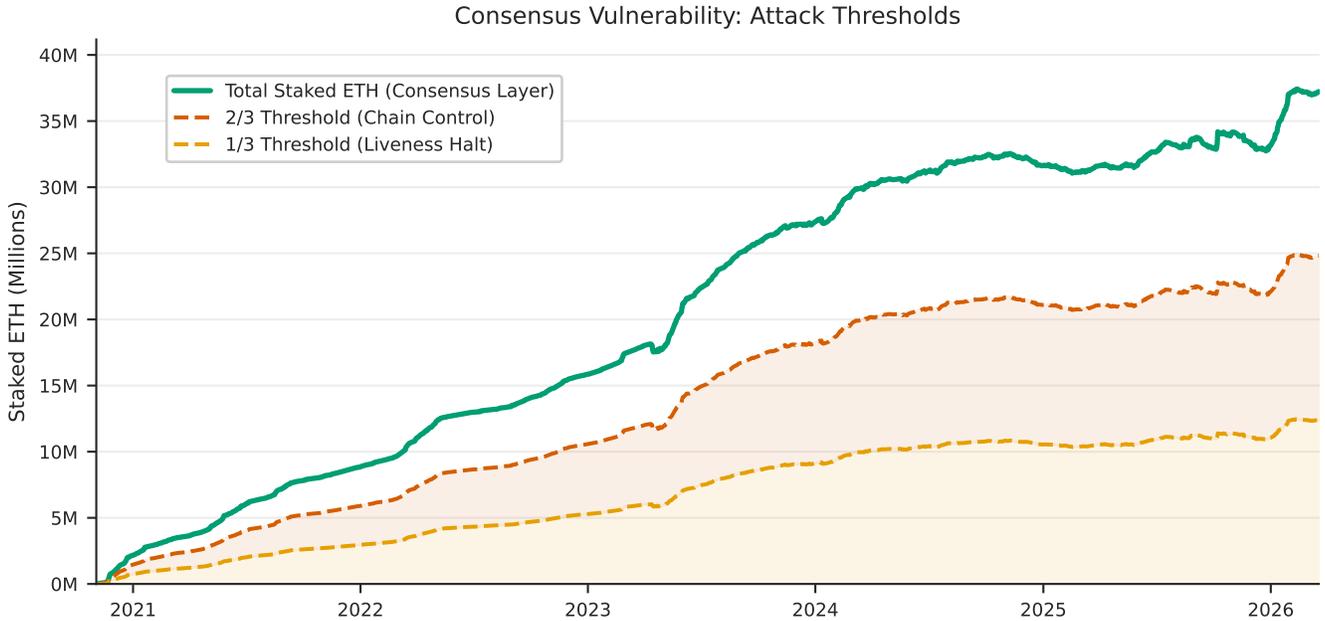


Figure 12. Historical accumulation of ETH in the Beacon Chain deposit contract (0x0...05fa). The green line represents the total stake securing the network (which relies on the vulnerable BLS12-381 curve for signature aggregation). The dashed lines identify the thresholds for quantum attacks: $> 1/3$ of the stake (yellow) allows an attacker to halt finality (liveness failure), $> 2/3$ (red) grants the ability to rewrite history and censor transactions (safety failure). Plot generated using data from `bigquery-public-data.crypto-ethereum` [195].

that this safety margin assumes a uniformly decentralized validator set, if stake remains concentrated in a few large pools (such as Lido accounting for approx. 20% of staked ETH [208]) with admin vulnerabilities, an attacker could target the specific key-management infrastructure of those providers to acquire a supermajority much faster.

Long-term remedy to the Consensus Vulnerability requires a phased transition to new validator credentials based on a suitable post-quantum multi-signature scheme. In order to help identify such a scheme, Ethereum Foundation has conducted research on candidate hash-based replacements for BLS12-381 [209]. In the meantime, a fast and effective mechanism for rotating validator keys could serve as a relatively simple stopgap measure for fortifying the Beacon Chain against earlier-than-expected quantum attacks on consensus. Following implementation and deployment of a post-quantum replacement and once a sufficient supermajority of stake has rotated to PQC validator credentials, the system may need to begin ejecting or slashing non-migrated validators to eliminate the lingering attack surface of the legacy system.

E. Data Availability Vulnerability

The Layer 2 ecosystem depends on the ability to commit bulk transaction data to the Layer 1 Ethereum blockchain. In order to check integrity of this data, validators used to have to download it via smart contract calldata which was a source of inefficiency. In order to address this throughput bottleneck, Ethereum introduced inexpensive, transient onchain data storage via 128 KiB blobs [210] together with the Data Availability Sampling (DAS) mechanism [173] which allows validators to probabilistically verify blob integrity by downloading only small, randomized data chunks called cells. Blobs and DAS led to massive speedups and lower costs for Layer 2 bulk transaction commitments.

DAS employs the KZG polynomial commitment scheme on the BLS12-381 elliptic curve. The binding property of the scheme is vulnerable to quantum attacks, allowing a quantum-capable adversary to forge cell authentication proofs and to deceive validators into believing that a blob is available despite it not being reconstructable from the cells with forged proofs. This can be used to stall rollups, creating opportunities for ransom attacks.

Furthermore, an attacker does not need to run Shor’s algorithm for each individual blob. The KZG scheme relies on the so-called trusted setup — a one-time cryptographic ceremony [211] to generate a secret random scalar encrypted in a collection of points on the BLS12-381 elliptic curve. The points comprise the scheme’s public parameters, known as the Structured Reference String (SRS), while the random scalar, sometimes referred to as “toxic waste”, must be

destroyed — a failure to erase it creates a permanent backdoor into the protocol. Mathematically, the toxic waste can be recovered as the discrete logarithm of two elliptic curve points in the SRS, so a CRQC can compute it from publicly available parameters. Once this secret is extracted, the adversary obtains a persistent, universal backdoor allowing them to forge data availability proofs at minimal computational cost and with no further need for a CRQC. In effect, a single successful quantum attack on the SRS creates a potentially tradable exploit that gives any adversary without a CRQC a persistent ability to attack DAS. The exploit remains functional until a new trusted setup is created. This constitutes an example of what we refer to as on-setup attacks.

Attacks against this single point of failure would have severe consequences for the L2 ecosystem. L2 rollups rely on L1 data availability to reconstruct their global state, so by withholding transaction batch data, an attacker effectively halts the L2 sequencers and validators. Throughput constraints and lower cost incentives have driven the majority of the ecosystem to this lower cost quantum-vulnerable architecture since the Dencun network upgrade in March 2024. Hence, the majority of modern Layer 2 data volume is now fundamentally dependent on the KZG scheme. As of February 2026, we estimate Data Availability Vulnerability risk is roughly 15 million ETH in TVS (Figure 11) which along with second order effects, such as complete liveness failure halting the affected L2 networks, would provide adversaries with leverage for ransom attacks against bridged L1 assets. As with Code Vulnerability, Stage 0 and Stage 1 Admins could partially mitigate impacts by pausing the Layer 2 activity until a workaround is found.

Mitigating the Data Availability Vulnerability requires upgrading the DAS protocol to a PQC commitment scheme before CRQCs come online. Anticipating this need, the Ethereum research community is exploring quantum-resistant alternatives, primarily hash-based polynomial commitments, such as Fast Reed-Solomon Interactive Oracle Proofs of Proximity [212], and Merkle-tree architectures [213]. While transitioning to these post-quantum alternatives will likely introduce higher bandwidth and storage overheads for validators compared to the highly efficient KZG commitments, replacing them is a mandatory step to ensure the long-term survivability of Ethereum’s L2 scaling infrastructure.

In summary, Ethereum has a broader overall quantum attack surface than Bitcoin. However, this is compensated by stronger community leadership in the Ethereum Foundation. Operating as a non-profit out of Switzerland since 2014, the Ethereum Foundation supports the ecosystem through funding, research and coordination [214]. In the spirit of decentralization, the Ethereum foundation does not “control the decision making process”, but can exert significant influence in advocating for the long-term success of the ecosystem by helping to organize development work, coordinate decision making, and define protocol strategy and roadmap. This influence was demonstrated when it successfully implemented a bailout of The DAO by carrying out a one-off irregular change of blockchain state to reverse the July 2016 hack [215–217]. With time, acceptance of emergency measures appears to have declined as a similar intervention in April 2018 was rejected [218]. Nevertheless, stronger leadership will likely make it easier and faster for Ethereum to transition to PQC [219] and to adopt any mitigation measures deemed necessary. Moreover, the proven willingness to reach for extraordinary means to preserve the integrity of the system provides a high degree of assurance. By contrast, Bitcoin’s decentralized community and the lack of a singular center of power may require a potentially drawn-out process of consensus building. Indeed, the increased storage and compute requirements of PQC may place the Bitcoin community in a position similar to August 2017 when divisions over block size change led to a hard fork that created Bitcoin Cash.

VI. QUANTUM VULNERABILITIES OF OTHER BLOCKCHAINS

Bitcoin and Ethereum dominate the cryptocurrency world in terms of economic activity, but a lot of financial and cryptographic innovation and experimentation takes place in the broader ecosystem and on other blockchains. Some of these developments seek to address quantum vulnerabilities while others continue to introduce new ones by building innovative features based on the ECDLP. The peril of the latter developments has been recognized in the community for example by Luke Parker, a prominent developer in the Monero ecosystem, who has called for a moratorium on research and development of quantum-vulnerable protocols [220, 221]. In this section, we briefly examine a broad sample of blockchains and cryptocurrencies from the standpoint of quantum security.

A. Bitcoin and Ethereum Derivatives

Among the many early cryptocurrencies based on Bitcoin, one can identify two distinct types. Codebase forks, such as Litecoin (LTC), Zcash (ZEC), and Dogecoin (DOGE), employ modified Bitcoin software and produce their own blockchain starting with a new genesis block. Chain forks, such as Bitcoin Cash (BCH) and eCash (XEC), run modified Bitcoin software and share a part of transaction history with their parent chain up to the block after which they split off as a hard fork. Many elements of our Bitcoin analysis apply to both types of forks with suitable adaptations to account for the software customizations, new features, and blockchain history.

For example, Litecoin implemented both SegWit and Taproot upgrades and our earlier discussion of these upgrades applies to it. In particular, litecoin secured by a P2TR locking script exposes the public key and is therefore vulnerable to at-rest attacks while litecoin secured by a P2WPKH script hides the public key behind a hash, so if the key is not reused or otherwise exposed elsewhere, then it is only vulnerable to on-spend attacks. At the same time, as we remarked earlier (see Figure 6), Litecoin’s 2.5 minute average block interval significantly reduces the probability of a successful on-spend attack using early fast-clock CRQCs. Nevertheless, theoretical advances and increases in the number of qubits are very likely to bring on-spend vulnerability to Litecoin eventually.

Dogecoin also eschewed the SegWit and Taproot upgrades and therefore does not support Bitcoin’s P2WPKH, P2WSH, and P2TR script types. Further, its one minute average block interval makes on-spend attacks using early CRQCs practically impossible under our current assumptions (see Figure 6). A recent draft Dogecoin Improvement Proposal [222] seeks to add native support for verification logic of quantum-vulnerable zkSNARKs with the objective of enabling L2 scaling solutions. Depending on the chosen ZK protocols, this may introduce a vulnerability to on-setup attacks to Dogecoin.

Bitcoin Cash launched in 2017 as a hard fork from Bitcoin due to disagreement about the solution to scalability issues. Bitcoin Cash never introduced SegWit or Taproot addresses and, like Zcash and Dogecoin, lacks the P2TR vulnerability. At the same time, as a hard fork of Bitcoin, it shares Bitcoin’s early transaction history and thus, inherits the problem of P2PK coins, including Satoshi era mining rewards.

Pegged sidechains [223, 224] are blockchains built in order to add new functionality to an existing network without creating a new store of value. They rely on bridges to transfer assets between the parent blockchain and the sidechain. For example, Rootstock (RSK) [225, 226] is an Ethereum-based Bitcoin sidechain with EVM-compatible stateful smart contracts and two types of accounts (EOAs and contract accounts) whose native token (RBTC) is pegged 1:1 to bitcoin. Users can send bitcoin to a special multi-signature Bitcoin address where it is received by an automated mechanism called PowPeg causing release of the corresponding amount of RBTC on the Rootstock sidechain where users can take advantage of a rich ecosystem of DeFi applications enabled by the EVM-compatible smart contracts. Later, they can send their RBTC to a precompiled smart contract called Bridge to trigger the release of the corresponding amount of bitcoin on the Bitcoin blockchain. In addition to the quantum vulnerabilities inherited from Ethereum, such as the Account and Code vulnerabilities, Rootstock introduces new quantum dangers. The system protects the private keys that control the multi-signature address using Hardware Security Modules (HSM). However, an attacker equipped with a CRQC can derive these keys without compromising the HSMs and steal any bitcoin residing at the address. This creates the risk of sidechain bankruptcy by making it unable to satisfy redemptions. Besides technical quantum vulnerabilities, Rootstock enables practices that may expose users to quantum risk. For increased convenience it allows users to omit the explicit Rootstock destination address from the deposit transaction on the Bitcoin side. In this case, the address is automatically derived as a cryptographic hash of the user’s public key. Consequently, the destination account on Rootstock ends up controlled by the same private key that signed the first input of the deposit transaction. This is convenient, but gives rise to Offchain Exposure vulnerability by exposing the bitcoin public key through the user’s activities on the Rootstock sidechain.

B. Privacy-Preserving Blockchains

In 2022, Litecoin implemented a private transactions protocol called Mumblewimble [227, 228] as a sidechain with blocks embedded in the canonical Litecoin blockchain. Inputs, outputs and amounts of Mumblewimble transactions are hidden or erased by a combination of Pedersen commitments, transaction aggregation, and pruning of spent outputs. In order to eliminate the need for interaction between the sender and receiver transacting on Mumblewimble, Litecoin introduced stealth addresses [167] and modified the protocol to use ECDH key exchange for offline derivation of secrets between the sender and receiver. Pedersen commitments and ECDH key exchange protocol are both vulnerable to quantum attacks. Moreover, the two elliptic curve points that Mumblewimble uses in the Pedersen commitments are fixed public parameters creating a vulnerability to on-setup attacks. Indeed, an adversary with a CRQC can solve the corresponding ECDLP once to manufacture a tradable exploit that can be used repeatedly to break the binding property of the Pedersen commitments and execute undetected inflation attacks without further need for a CRQC.

The Litecoin community considered a variety of protocols for private transactions, including post-quantum zk-STARKs. However, they ultimately rejected this option due to its high resource cost [228]. The design provides the option to switch from Pedersen commitments to ElGamal commitments to protect the monetary base at the expense of privacy [228, 229]. ElGamal commitments are quantum vulnerable, but in a different way than Pedersen commitments. Before the switch, a quantum adversary can steal and create new coins, because the binding property of Pedersen commitments is vulnerable to quantum attacks. In addition, a quantum adversary can learn transaction amounts, despite quantum resistance of the hiding property of Pedersen commitments, by exploiting quantum vulnerability in ECDH key exchange to learn blinding factors. After the switch, a quantum adversary can learn transaction

amounts due to quantum vulnerability in the hiding property of ElGamal commitments, but cannot steal or create new coins, because the scheme’s binding property is quantum-resistant.

Zcash (ZEC) is a Bitcoin-based blockchain that employs advanced cryptography to provide confidential transactions. Zcash wallets use zkSNARKs to prove to the network that a transaction is valid without revealing the sender, recipient, or amount sent. Zcash has evolved over three generations of protocols for shielded transactions [165]. Sprout, the original shielded protocol based on the BCTV14 proof system [230], was launched in 2016 and phased out in 2020 [231, 232]. Sapling, launched in 2018, implemented the Groth16 protocol [233] on the BLS12-381 curve, substantially improving performance. The launch was preceded by a trusted setup ceremony [234] which produced public protocol parameters called the Structured Reference String (SRS). As in the case of Ethereum’s DAS, the Sapling ceremony also generated a secret scalar, known as the “toxic waste”, which had to be destroyed in order to avoid leaving behind a permanent backdoor into the protocol. The toxic waste can be recovered from the SRS using a CRQC to create a persistent tradable exploit enabling non-quantum attackers to stealthily inflate monetary base in the Sapling shielded pool. The need for trusted setup and the vulnerability to these on-setup attacks were both eliminated in the newest shielded protocol called Orchard that launched in 2021. The protocol introduced the Halo 2 scheme [235] using Pasta elliptic curves [181] and increased scalability by introducing recursive proof composition. Nevertheless, Orchard’s improved quantum security posture does not eliminate all quantum risks.

Indeed, many Zcash innovations, including those in the Orchard protocol, use ECDLP-based cryptographic primitives, such as zkSNARKs, Pedersen commitments, and ECDH key exchange, and are therefore vulnerable to quantum attacks. For example, Zcash enables users to publish an effectively unlimited number of public addresses, known as diversified addresses, associated with a single private incoming viewing key. A classical adversary cannot determine whether different diversified addresses belong to the same user — a privacy feature known as unlinkability. However, the incoming viewing key is the solution to the ECDLP associated with the diversified address, so a quantum attacker can defeat unlinkability by deriving the incoming viewing key. Practical quantum attacks on unlinkability would likely be done offline. Potential online attacks on diversified addresses would not benefit from the factor-of-two speed-up achieved by “priming” the quantum computer which we described earlier, because priming depends on the knowledge of the generator which is part of a diversified address, not a fixed constant publicly known before the attack. Consequences of quantum attacks on diversified addresses are limited to the recovery of the incoming view key by Zcash’s robust key hierarchy which denies quantum attackers access to the user’s spending key.

Another quantum vulnerability concerns encryption. Even though Zcash notes are encrypted onchain using a quantum-resistant authenticated symmetric encryption scheme [236], the encryption key is derived using ECDH from the recipient’s incoming viewing key and an ephemeral private key. The ephemeral public key is recorded on the blockchain alongside the encrypted note. Consequently, a quantum attacker who knows the target diversified address can recover the encryption key and decrypt the note, including transaction amount and memo. Thus, the most pressing quantum danger for Zcash is the eventual retroactive degradation of privacy by future quantum attacks on known addresses. Beyond privacy, CRQCs will be able to compromise soundness of zkSNARK protocols raising the risk of theft and creation of counterfeit notes. The Zcash community is discussing proposals for addressing quantum vulnerabilities, such as recently proposed foundational features to enable recoverability from quantum attacks [237]. This feature is a part of Zcash’s broader plans for post-quantum transition [238].

Each of Zcash’s shielded transaction protocols — Sprout, Sapling and Orchard — has its own distinct value pool with the fourth pool for transparent transactions inherited from Bitcoin. Assets moving between different shielded pools must pass through the transparent pool which allows Zcash to track total monetary supply within each shielded pool using a mechanism called Turnstile [239, 240]. This provides the last line of defense against supply inflating attacks [238]. Zcash never implemented Bitcoin’s SegWit and Taproot upgrades, so transparent transactions use P2PKH and P2SH script types. On-spend attacks against Zcash are made extremely challenging under our current assumptions by its 75 seconds average block interval (see Figure 6).

C. Post-Quantum Blockchains

A few blockchains have made progress in real-world deployment of PQC. In particular, the QRL [63, 64] launched in 2018 stands out as post-quantum from inception. Its original design was based on the stateful post-quantum signature scheme known as XMSS [241] and it is currently adding support for the stateless post-quantum signature scheme called CRYSTALS-Dilithium [242] and recently standardized by NIST under the name ML-DSA [243]. Other examples of post-quantum blockchains include Mochimo (MCM), which uses a variant of hash-based post-quantum Winternitz One-Time Signatures (WOTS) [244, 245], and the post-quantum privacy-preserving Abelian blockchain (ABEL), which makes extensive use of lattice-based PQC.

Algorand (ALGO) provides an example of real-world deployment of PQC on an otherwise quantum-vulnerable blockchain. It launched in 2019 as a Pure-Proof-of-Stake blockchain for smart contracts and fast transactions. Smart

contracts on Algorand are written in popular high-level programming languages like Python and TypeScript and are compiled to assembly-like Transaction Execution Approval Language (TEAL) [246] which executes on the Algorand Virtual Machine (AVM) [247]. In addition to builtin single- and multi-signature transactions, Algorand supports stateless smart signatures and stateful smart contracts for DeFi applications. Algorand’s consensus and builtin transactions are based on quantum-vulnerable Ed25519 digital signature scheme. However, it has recently deployed post-quantum Falcon digital signatures [248, 249] for smart transactions and state proofs (cryptographic attestations of blockchain state for cross-chain integrations). Algorand has also made Falcon signature verification available as a TEAL primitive [250] to enable development of quantum-safe smart contracts for AVM. These PQC technologies are now publicly available: Algorand executed its first PQC-secured transaction in 2025 [62]. Moreover, Algorand enables users to change the private keys associated with their accounts [251]. While this mechanism does not provide full quantum security at present, it facilitates future PQC migration.

There are also experimental and test deployments of PQC on quantum-vulnerable blockchains. For example, Solana (SOL) [52] deployed an experimental feature called Solana Winternitz Vault [252, 253] which uses WOTS to protect digital assets. More recently, the XRP Ledger (XRP) [56, 57], deployed post-quantum ML-DSA signatures on its AlphaNet test instance [67]. The blockchain provides extensive support for RWA tokenization including compliance controls, issuer permissions and asset metadata [60] and is increasingly utilized by global financial institutions. It currently holds about two thirds of all TBILL tokens backed 1:1 by short-dated U.S. Treasury bills [254] with most of the remaining one third held on Ethereum.

D. Stablecoins and Real-World Asset Tokenization

Involvement of global financial institutions and RWA tokenization reflect the significant growth and diversification in the digital economy and led to the introduction of new types of digital assets. A prominent new asset class is known as stablecoins which are digital tokens that seek to maintain a stable exchange rate to a fiat currency, commodity or another cryptocurrency, either by means of algorithmic trading or maintenance of liquid reserves. The latter type of stablecoins are a prime example of tokenization where the RWA backing the digital token is a fiat currency.

Unlike Bitcoin and Ethereum, many new cryptocurrencies, called crypto tokens, do not have their own dedicated blockchain infrastructure and are instead implemented as smart contracts on other blockchains, primarily Ethereum. In particular, the two most popular stablecoins, Tether (USDT) and USD Coin (USDC), are both available as ERC-20 tokens on Ethereum. The smart contracts that govern these cryptocurrencies give special administrative privileges to a handful of accounts. These include the right to mint and burn coins, freeze accounts and, most importantly, to upgrade the contract’s logic. In particular, minting of unbacked tokens could potentially cause a market-driven peg collapse. It may be possible to contain some of the damage caused by such an attack by early detection and halt in trading. However, the mechanisms involved may need to involve offchain elements, because the logic upgrade privilege confers essentially unlimited control over the onchain contract. Access to this privilege is typically secured by a multi-signature that requires a certain threshold number of valid digital signatures from a group of accounts controlled by important stakeholders. Due to past upgrades, these public keys are exposed on the blockchain which creates the existential risk of a quantum attacker taking complete control over the smart contract. This illustrates the fact that crypto tokens, such as USDT and USDC, are exposed to highly concentrated quantum risk arising from the Admin Vulnerability in the host blockchain.

Cross-chain assets, such as USDT and USDC, may have quantum vulnerabilities on many blockchains and may inherit quantum weaknesses from interoperability mechanisms, such as bridges and Interblockchain Communication Protocol [255]. However, their presence on multiple blockchains also introduces a potential new element for the mitigation strategy. As some blockchains adopt PQC and others lag, multi-chain cryptocurrencies will have the option to withdraw from quantum-vulnerable systems while relying on post-quantum blockchains to continue their economic activity securely. There is precedent for a tenant cryptocurrency withdrawing from a host blockchain. For example, in February 2024, Circle announced [256] that they ceased minting new USDC on the TRON blockchain and allowed users to transfer their coin to other chains until February 2025. The one-year gap between the cessation of minting and termination of redemptions highlights the potential risks and liabilities associated with the long-tail of lingering assets and indicates the need for forward-looking assessment of a host blockchain’s vulnerabilities. The possibility of cross-chain migration of high profile tenant cryptocurrencies to post-quantum systems may exert competitive pressure on blockchains such as Ethereum to accelerate PQC adoption. In fact, early developments in this direction have already taken place, for example USDC is available on the Algorand blockchain which supports post-quantum digital signatures. The competitive pressure on the Ethereum ecosystem may further increase with the emergence of EVM-compatible post-quantum blockchains. In fact, the post-quantum Abelian blockchain already supports EVM-compatible smart contracts via a zk-rollup called QDay. Furthermore, QRL’s development roadmap [257] lists the support for EVM smart contracts as one of its key objectives.

E. A Taxonomy of Quantum Risk Profiles for Distributed Ledgers

The variety of quantum risk profiles of popular blockchains can be summarized using four coarse categories. The first one includes post-quantum blockchains, such as the QRL, Mochimo, Abelian and QDay (Abelian’s L2 network).

The second one consists of protocols in which it is possible for individual users to avoid long-term exposure of quantum-vulnerable public keys. This category includes UTXO-based ledgers, such as Bitcoin, Litecoin, Dogecoin as well as Cardano [130] which uses UTXOs extended with logic and data to support smart contracts [258]. These protocols do not tie wealth to static account identities and allow users to secure assets with ephemeral public keys hidden behind a cryptographic hash. This makes it possible for cautious users to evade at-rest attacks leaving only the possibility of on-spend attacks that are more challenging for a quantum attacker. In fact, in our Scenario 2, early CRQCs are too slow to launch these on-spend attacks. However, the potential resistance of these blockchains to at-rest quantum attacks is often lost in practice due at-rest vulnerabilities arising outside the basic transaction protocol. For Bitcoin and its derivatives, they primarily stem from public key exposure driven by user practices such as address reuse, whereas in Cardano, at-rest vulnerabilities mainly affect staking activity in its Proof-of-Stake consensus protocol as well as voting in its decentralized governance system and originate from onchain exposure of staking and voting keys (distinct from spending keys) in registration certificates [259]. Thus, an individual user’s ability to avoid public key exposure does not preclude systemic risks, such as those arising from vulnerabilities in consensus and governance.

The third category consists of protocols which make long-term exposure of quantum-vulnerable public keys inevitable. It includes blockchains with persistent accounts, such as Ethereum, Solana, Rootstock, Algorand, TRON and the XRP Ledger. These blockchains utilize an account model and either use public keys directly as account addresses (Solana) or expose them in the first transaction. Simultaneously, their ecosystems increasingly leverage on-chain account histories to assess financial risk and confer privileges making account migrations costly for users. This design makes active public keys easier for quantum attackers to find compared to UTXO-based blockchains. Moreover, for legacy accounts on Ethereum and Rootstock, and for standard keypair accounts in Solana, this is exacerbated by lack of support for key rotation which effectively locks users into public key exposure. Modern Ethereum, Solana and Rootstock accounts are controlled by smart wallets and support key rotation, but legacy accounts remain a lingering vulnerability. By contrast, Algorand [251], TRON [260] and the XRP Ledger [261] support native, protocol-level key rotation. Public blockchains in this category generally support smart contracts. Consequently, the scale of funds at stake in quantum attacks on these blockchains is greatly increased by smart contract vulnerabilities and is expected to grow further due to rising popularity of stablecoins and tokenization. Indeed, potential compromise of administrative keys for a high-value bridge may represent a systemic existential threat.

Finally, the fourth category comprises the quantum vulnerable privacy-preserving blockchains, such as Zcash, Monero and Litecoin’s Mumblewimble sidechain. In addition to the forward-looking threats, such as asset theft and supply inflation, the privacy-preserving blockchains face retroactive degradation of privacy. An adversary with a CRQC could potentially deanonymize years of historical confidential transactions for known addresses, making PQC migration of these blockchains not only a way of preventing future quantum attacks, but also a means of reestablishing privacy today and thus a critical imperative.

VII. RISKS AND CHALLENGES IN MIGRATING TO POST-QUANTUM CRYPTOGRAPHY

The root cause of the numerous quantum vulnerabilities discussed in this whitepaper lies in the widespread use of cryptographic protocols, especially digital signatures, based on ECDLP — a computational problem believed to be hard for classical computers that is known to be efficiently solvable on quantum machines [262]. Therefore, the only durable long-term solution to these vulnerabilities lies in upgrading the underlying schemes to post-quantum alternatives [263].

The task of upgrading blockchains to PQC faces significant headwinds in multiple areas. With few exceptions [264], PQC protocols are relatively new and have faced less scrutiny and seen less real-world usage than older protocols based on ECDLP. They are divided into five categories [265] depending on the choice of the computational problem suspected to be hard for classical and quantum computers, that they use to replace ECDLP as a hardness assumption. The categories include cryptosystems based on lattices [266], hash functions [267], codes [268], multivariate polynomials [269], and isogenies [270].

The relative novelty of these schemes raises justifiable concerns about the possibility of undiscovered weaknesses against both classical and quantum attacks. Indeed, the Supersingular Isogeny Diffie-Hellman protocol [271] has already succumbed to a classical attack [272]. Security of other post-quantum cryptosystems is an active area of research. In particular, research on quantum algorithms based on Regev’s reduction [273] has recently led to the discovery of a novel quantum algorithm called Decoded Quantum Interferometry (DQI) [274]. Even though the

algorithm tackles an approximate optimization problem rather than a cryptographic task, algorithms based on Regev’s reduction are intimately connected with the Short Integer Solution (SIS) [275] and Learning With Errors (LWE) [273] problems which feature prominently in hardness assumptions of some of the lattice-based post-quantum cryptosystems. A different line of research, investigating quantum algorithms based on the Kikuchi method [276, 277], yielded a new quantum algorithm [278] that achieves a quartic (square of square) speedup over the best known classical algorithm for sparse Learning Parity with Noise (LPN), a problem related to LWE. It is conceivable that further progress in these active areas of research yields unexpected breakthroughs regarding quantum attacks on lattice-based cryptosystems. On the other hand, if these lines of research yield no efficient algorithms for SIS or LWE, then the extra scrutiny will further boost confidence in lattice-based cryptography. Currently, lattice- and hash-based cryptosystems appear to be the most promising of the PQC schemes and have been standardized by the U.S. National Institute for Standards and Technology (NIST) [241, 243, 279, 280].

A further complication arises from additional requirements of blockchain applications, such as signature aggregation and recursive proof aggregation. Such schemes are needed to aggregate large numbers of signatures in Ethereum’s Proof-of-Stake consensus mechanism and would enable blockchains to maintain high transaction rates. Ethereum Foundation has supported and conducted research on analyzing suitable PQC protocols, such as certain variants of the hash-based XMSS signature scheme [209].

Software implementations of PQC schemes are another source of uncertainty. They are inevitably newer and less battle-hardened than older libraries for ECDLP-based schemes, so any upgrade that introduces them into a blockchain software stack brings with it a chance of new software bugs and subtle security weaknesses. For example, some LWE-based schemes, such as Falcon [248, 249], involve sampling from discrete Gaussian distribution which has in the past created side-channel vulnerabilities [281].

Use of post-quantum signature schemes also increases resource requirements. Indeed, a key reason for the choice of ECDLP-based signature schemes in blockchain networks is their excellent performance: the popular variants of these signature schemes use small public keys and signatures, which reduces storage and bandwidth requirements, and are efficient to sign and verify, which reduces compute requirements. Unfortunately, none of the PQC signature schemes matches these excellent performance characteristics. Indeed, the storage and compute requirements of PQC schemes are typically one or more orders of magnitude greater than those of the legacy ECDLP-based protocols [25, 282]. Indeed, the high cost of verifying SQISign signatures [283] raised concerns about the possibility of denial of service attacks [284]. Storage requirements of post-quantum Falcon signatures used on Algorand blockchain are about 1280 bytes in size while quantum-vulnerable ECDSA signatures used in Bitcoin are 64-73 bytes. These increases also affect usability as they may lead to much larger payment addresses [285].

Moreover, a common approach to fortifying digital signatures against possible weaknesses in a novel post-quantum scheme is to use composite signatures [236, 286] that combine a proven scheme, e.g., based on ECDLP, to protect against forgery by classical computers and a novel PQC scheme, e.g., based on LWE, to protect against forgery by quantum computers. This choice enhances security at an even greater performance cost than the PQC scheme alone.

If Bitcoin keeps the block size unchanged while introducing post-quantum signatures, then fewer transactions will fit into a block, reducing future transaction rate. However, increasing the block size proved controversial in the past. In particular, the August 2017 hard fork that created Bitcoin Cash highlights the divisiveness of issues surrounding design choices that affect compute resources needed to run a Bitcoin full node. These choices raise fears that increased resource requirements portend network centralization, which is at odds with some of the foundational ethos of Bitcoin.

The resource costs of post-quantum digital signatures could be absorbed by L2 scaling solutions, amortized using signature aggregation [209, 287] or reduced via succinct proofs [288], such as post-quantum hash-based SNARKs, that allow recursive proof composition [289] and promise high efficiency.

Once PQC is adopted, digital asset holders will generally need to initiate transactions to send their assets from old quantum-vulnerable addresses to new post-quantum ones. On some blockchains, this process is expected to run into technical bottlenecks. On the Bitcoin blockchain, at the current transaction rate, the migration would take several months even if the network processed only asset migrations [28]. Thus, ideally, the process should start years before the quantum threat truly emerges in order to avoid causing problematic network congestion. We expect this asset migration to take less time on younger blockchains and on blockchains that implement the account model, such as Ethereum, due to the resulting asset consolidation. The dangers associated with the long tail of asset migration can be mitigated using commit-reveal schemes [290] which can be used to protect existing script types against on-spend attacks at the cost of increased complexity and duration of transactions involving non-migrated assets.

A related issue concerns the need to alert and educate a very large number of cryptocurrency holders about the need to migrate to post-quantum addresses. This calls for updates to User Interfaces (UIs), ensuring quantum-safe default settings, and possibly even coordinated messaging campaigns.

The technical and social complexities of switching blockchains to post-quantum signature schemes indicate that the process will take years and cannot be delayed until the exact timeline and feasibility of constructing CRQCs become completely clear (and again, those details might not be broadcast publicly). At the same time, these complexities

and challenges are feasible to overcome as demonstrated by Algorand, Solana and the XRP Ledger which have made notable progress in real-world adoption of PQC, as well as by QRL, Abelian and Mochimo which have used PQC protocols from inception.

Nevertheless, even when blockchains implement post-quantum signature schemes and asset owners are able to transact out of quantum-vulnerable to quantum-safe addresses, many blockchains will still contain quantum-vulnerable assets, such as Bitcoin’s P2PK coins. The P2PK locking scripts have seen little use since 2009 and a majority of the 1.7 millions bitcoin have not been transacted since 2009 despite the fact that such coins are known to be vulnerable to at-rest quantum attacks. This leads to the widespread assumption that (in most cases), the private keys of these Bitcoin UTXOs have been lost. For example, accounts believed to have been owned by Satoshi Nakamoto hold about 1 million bitcoin [291, 292] behind these P2PK scripts and the most common theories are that Satoshi has either died or deliberately burned his private keys (or never recorded them in the first place) either out of some sort of ideological conviction or because they were almost worthless at the time they were initially mined [293–296]. If true, such assets can never be safely migrated to a more secure protocol and remain fixed on the ledger without those private keys. Thus, there is no obvious way to update Bitcoin in a fashion that would keep lost P2PK coins and other dormant assets spendable by their owners and also protected against quantum computers. This makes them an especially enticing target for quantum computers and pose a policy challenge for the Bitcoin community and governments.

VIII. DORMANT DIGITAL ASSETS

A. The Challenge Posed by Dormant Assets

Inevitably, some vulnerable assets will not migrate to post-quantum protocols in time or possibly ever, perhaps because their owners do not learn of the threat until it is too late or perhaps because they have lost their private keys. The Ethereum blockchain’s contract accounts present similar long-tail migration issues. Thus, in addition to planning and executing upgrades to cryptographic protocols, each cryptocurrency community also faces challenges regarding quantum-vulnerable assets and smart contracts that may linger on public blockchains for an extended or indefinite period of time.

Despite lack of unambiguous precedent, many jurisdictions could classify accessing abandoned cryptographic assets, such as the P2PK coins, without authorization as theft. However, we maintain that if protocol changes are not made, vulnerable assets will eventually be cracked by quantum computers and taken irrespective of the law. In the absence of a clear resolution, these assets are likely to become a lucrative target for bad actors. We quantify the scale of some of the dormant assets at stake in Figure 13.

Quantum computing providers could put in place systems to detect and deny workloads involving Shor’s algorithm, but sufficiently sophisticated actors may find ways to circumvent such safeguards, e.g., by using circuit obfuscation schemes [298, 299]. Dormant digital assets, such as P2PK coins, are a sufficiently high value target that the asymmetry of cyberdefense suggests an “assume breach” mindset. If taking the coins is clearly deemed illegal, then in the absence of further policy changes, possible outcomes include: A technologically sophisticated adversary builds their own CRQC and uses it to steal vulnerable bitcoins. An adversary hacks into a CRQC and uses it to steal vulnerable bitcoins. An adversary coerces (e.g., through extortion, bribery or violence) people known to have access to a CRQC into stealing or facilitating the theft of vulnerable bitcoins. A customer of a company cleared to use the device misappropriates their access to steal vulnerable bitcoins or their access is commandeered by a hacker. This list is not exhaustive. We cannot predict exactly how these coins would be taken but we maintain that if Bitcoin protocol changes are not made, they will eventually be taken (perhaps using devices based in the U.S., perhaps not). Without protocol changes, laws criminalizing the taking of these coins merely ensure they are eventually taken by nefarious agents. This possibility, combined with the enormous value of the assets in question, calls for a resolute policy response.

In the rest of this section we discuss two complementary sets of policy options available to address the challenge of dormant assets. First, we consider the three main policy options discussed in the Bitcoin community, called Do Nothing, Burn, and Hourglass. We also propose a solution based on a new dedicated sidechain that employs offchain cryptographic proofs of ownership to maximize the amount of dormant assets that return to their rightful owners. We call our approach “bad sidechain” since it is inspired by the notion of “bad bank” [300] used for orderly resolution of distressed assets in traditional finance. Next, we consider public policy responses available to governments. We begin by explaining the practical reasons why legally mandating the destruction of the dormant assets is not viable. Finally, we describe three realistic public policy options: regulated digital salvage, national security response, and government engagement with the Bitcoin community.

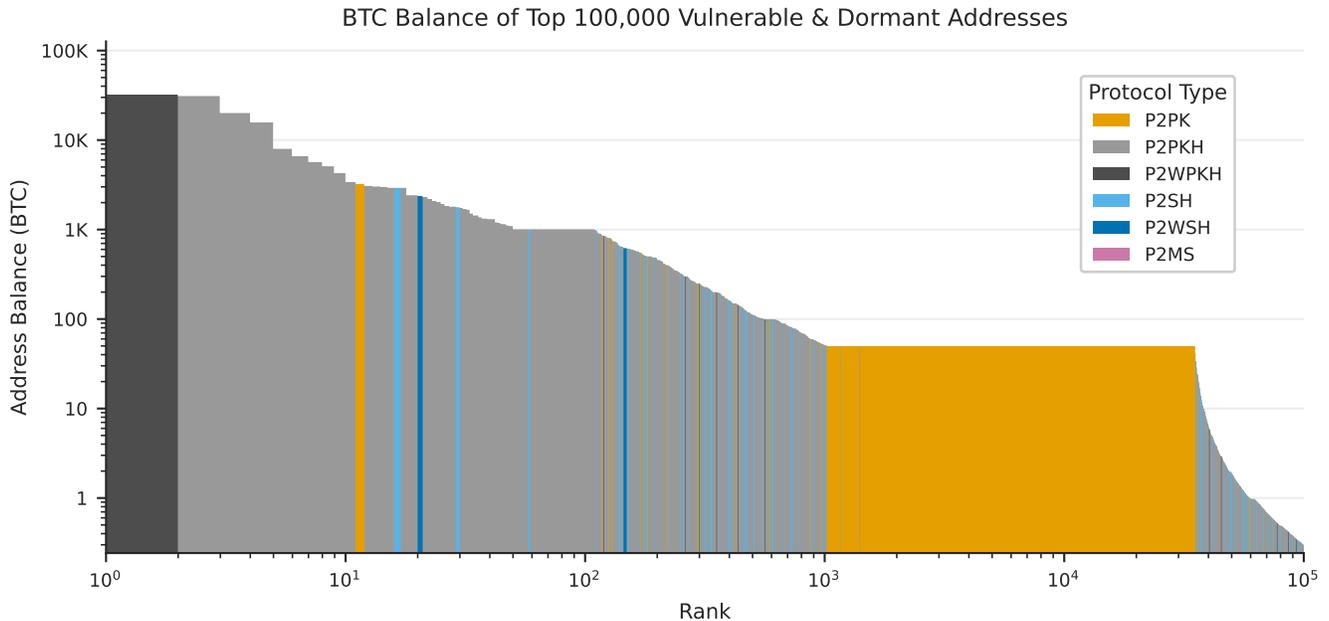


Figure 13. Similar to Figure 7, this plot displays the BTC balance of the top 100,000 Bitcoin addresses ranked by value that are vulnerable to at-rest attacks from exposed or re-used keys and have not initiated a spend in the last 5 years (This ignores things like receiving dust [297]). The sum of all BTC in vulnerable addresses in this list is ~ 2.3 million BTC. Plot generated using data from `bigquery-public-data.crypto_bitcoin` [139].

B. Bitcoin Community’s Options for the Challenge of Dormant Assets

1. Do Nothing, Burn, and Hourglass

The Bitcoin community (especially those with large Bitcoin holdings) might not be enthusiastic to have quantum computers salvage the lost Satoshi era coins. Any sudden increase of total Bitcoin supply would likely lead to a decrease in the value of Bitcoin. We note that the extent of this depends strongly on how, and how quickly, the coins are reintroduced and how many new buyers step forward. Because large quantum salvage operations would themselves quickly become significant holders of the cryptocurrency, they would have incentive to take steps to safeguard its overall value, for example by reintroducing the coins slowly into circulation. We expect the total trove of P2PK coins to take months (in Scenario 1) or years (in Scenario 2) to unlock, see Figure 14.

The Bitcoin community is considering three general categories of approaches to dealing with dormant assets [26]: Do Nothing: the protocol is kept unchanged allowing quantum attackers to acquire the dormant assets. Burn: the protocol is modified to render dormant assets unspendable after a certain date, effectively confiscating all assets locked by P2PK scripts. Hourglass [301]: the protocol is modified to limit the rate at which dormant assets may be spent. These proposals apply a one-size-fits-all policy blind to potential offchain proofs of ownership and are therefore typically only proposed as a means of dealing with the narrow problem of P2PK coins. Later, we contribute a fourth type of solution that takes offchain proofs of ownership into account, applicable to other classes of dormant assets.

The Do Nothing approach is motivated by the desire to protect property rights, avoid setting a confiscatory precedent and preserve Bitcoin’s fixed monetary supply, but risks price volatility as the dormant assets are reintroduced into circulation. By contrast, Burn prevents potential supply shock, but expropriates dormant assets. Indeed, any bitcoin owner who does not transact into a post-quantum protocol in time would permanently lose their assets. Burn could also potentially hurt the price of Bitcoin due the damaging precedent it sets. Accordingly, the Burn proposal has proven controversial. Hourglass seeks to strike a balance between these competing concerns by avoiding confiscation while stabilizing Bitcoin price. It also sets up a bottleneck for spending dormant coins which may incentivize parties with CRQCs to enter into bidding wars that transfer some of the abandoned wealth to Bitcoin miners.

In their most straightforward incarnation, Burn and Hourglass proposals entail modifications to the Bitcoin protocol that impose restrictions on transactions spending dormant coins: either by banning them outright after a certain date (in the case of Burn) or limiting their number to for example one per block (in the case of Hourglass). The key challenge for these proposals arises from the need for a broad consensus necessary to implement protocol changes.

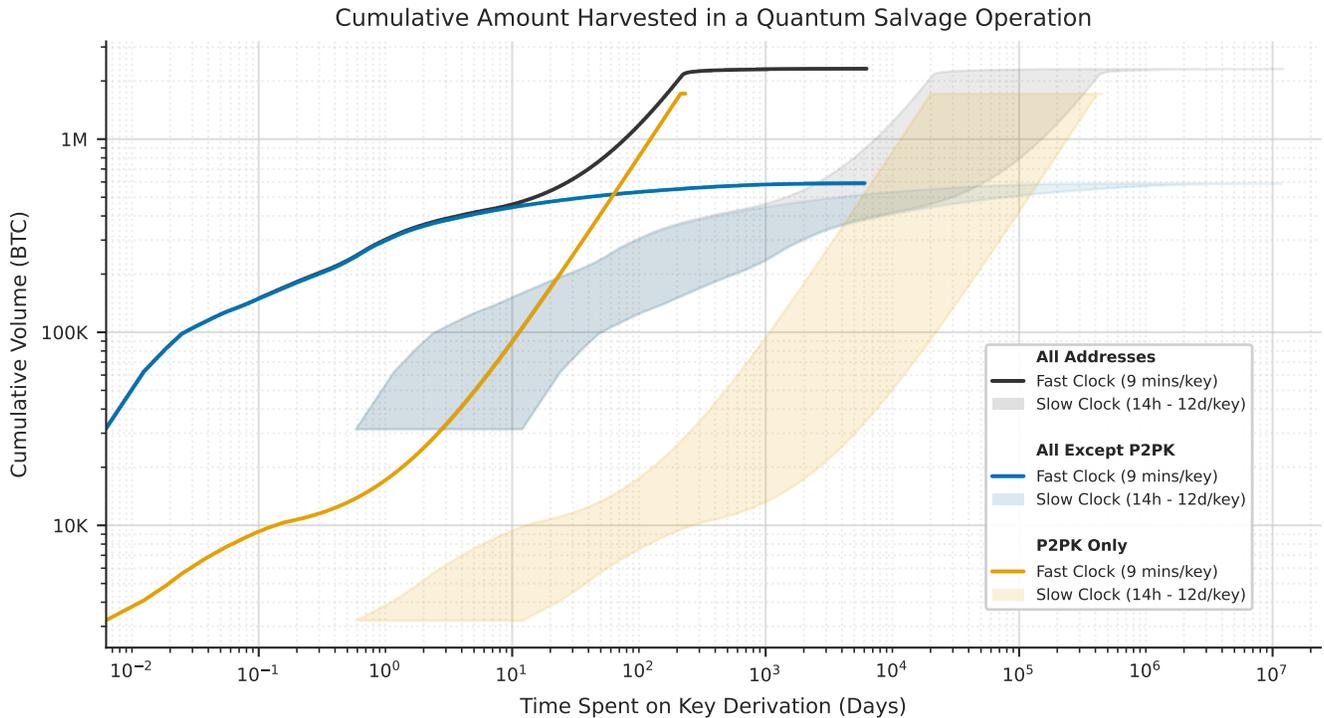


Figure 14. Cumulative amount of money harvested in a quantum salvage operation. We consider a single fast-clock quantum computer or a single slow-clock quantum computer deriving private keys for public keys exposed in P2PK addresses in decreasing order of amount locked. We assume 9 minutes per key derivation on the high clock speed architecture, where we are guided by our rigorous resource estimates, and between ~ 14 hours and ~ 12 days per derivation on the slow-clock speed architecture, where we employ a simplified model dominated by magic state cultivation with 10,000 to 200,000 qubits dedicated to cultivation (yielding 200 and 4000 magic T states per second, respectively). The first wallet holds approximately 2 thousand BTC while the total amount harvested is approximately 2.3 million BTC. Only public keys inactive for at least five years are included. Plot generated using data from `bigquery-public-data.crypto_bitcoin` [139] combined with chainstate database and raw blockchain data to account for the number of private keys an attacker must derive for each multi-signature address.

The Bitcoin protocol is maintained by a decentralized community of developers and the core network consists of miners (which validate transactions and bundle them into new blocks) and full nodes (which relay blocks and transactions and maintain a full and independent record of the entire blockchain). Software developers, businesses operating mining pools, bitcoin holders, and individual users running full nodes, who together form the Bitcoin community, need to work together to implement and adopt protocol changes, such as Burn and Hourglass.

There are two ways that technical changes can be made to the Bitcoin protocol. The most desirable type of update is known as a “soft fork”. A soft fork is possible when changes are “backwards compatible” and a decisive majority of miners and node operators agree to implement a proposal. The basic idea of a soft fork is that as long as the proposed change is narrowing the protocol rules (the new rules are a subset of the prior rules, e.g. banning certain UTXOs) then as long as a majority only acknowledge blocks that follow these rules, the longest blockchains (and thus, ultimately, the blocks that become permanent) will also follow these rules. However, if the change does not constitute a narrowing of the consensus rules or a clear majority cannot be established, then the only way to push a change is via hard fork. The community generally tries to avoid hard forks, because they result in multiple versions of the cryptocurrency, as happened for example in the 2017 hard fork that created Bitcoin Cash. From a technical perspective, both Burn and Hourglass proposals could be implemented as a soft fork since banning or limiting transactions that spend dormant coins is a narrowing of the protocol rules. However, this requires broad support in the community.

Consensus remains elusive. Indeed, different members of the Bitcoin community have different philosophical beliefs and financial incentives. Those who consider digital property rights fundamental tend to have strong objections to the Burn proposal. Large Bitcoin holders are likely concerned about a potential supply shock and its effect on Bitcoin price. Miners may welcome Do Nothing and Hourglass proposals due to potential increase in transaction fees and volumes. The diversity and complexity of the Bitcoin community makes the ultimate outcome of these ongoing debates hard to predict. Indeed, an informal poll at 2025 Presidio Bitcoin Quantum Summit in San Francisco saw

roughly equal support for each of the three categories of solutions [26]. In the event of a hard fork, every owner of the original digital asset becomes an owner of an equal amount of each of the two cryptocurrencies emerging from the fork. Relative value of the two assets depends on the ensuing market dynamics. Example developments that may affect relative price of the two cryptocurrencies include significant sales of one of the assets by a major cryptocurrency holder, especially when combined with simultaneous purchases of the other as well as new buyers stepping forward to build up positions in reaction to the fork’s effects on predictions of future returns. By contrast, miners generally follow the changing valuations of digital assets.

Irrespective of the outcome of the debate, governments will need to decide whether or not to legalize the salvage of dormant assets for the reasons outlined in the prior section. Indeed, it might be reasonable policy for governments to both legalize the salvage of dormant assets while also recommending that the Bitcoin community burn them. This is not contradictory because it would be taking a position that it is best for the Bitcoin community to protect itself from any adverse economic consequences of quantum salvage, while also recognizing that it cannot control how the protocol is updated.

2. *Bad Sidechain*

Even though the Bitcoin blockchain does not provide any mechanism to distinguish the true owner of dormant assets from someone who used a CRQC to derive the private key, for some assets there are offchain mechanisms to do this. Indeed, many public keys are derived from a mnemonic code — a group of easy to remember words — [302] which a quantum attacker cannot recover, but the original owner could present as proof of ownership. Similarly, users can create offchain cryptographic proofs of ownership in systems such as Project 11’s yellowpages registry [303].

None of these offchain proofs of ownership is accepted on the Bitcoin blockchain, but the community could launch a special-purpose pegged sidechain [223, 224] where CRQC operators could send the recovered dormant assets, such as P2PK coins, for resolution. The sidechain would implement manual or automated processes for identifying true owners of the original coins and would transfer the assets back to the Bitcoin network once the ownership question is resolved. The recovery sidechain could be derived from a post-quantum variant of Ethereum to enable the use of smart contracts to automate some or all of the resolution processes and its bridge to Bitcoin could be based on the Rootstock blockchain [225] which brings EVM-compatible smart contracts to Bitcoin. CRQC operators may be incentivized to send dormant coins to the recovery sidechain by contingent fees or forced to do so by a protocol upgrade, law, or both.

To expedite recovery, the setup could allow greater centralization than what is acceptable on the Bitcoin blockchain. For example, it could be governed by a consortium of stakeholders, perhaps organized in analogy to Liquid Federation that runs the Liquid Network sidechain [304]. The consortium could approve offchain ownership proof mechanisms, set recovery fees, decide burn schedules, adjust the peg and manage the pace of return of recovered coins to control inflationary effects. At a high level, the sidechain would serve a function bearing some analogy to the “bad bank” used in the resolution of distressed assets on the books of a failing bank in traditional finance [300].

In the presence of a broad agreement on what types of proofs of ownership are acceptable, one can imagine a decentralized market ecosystem of multiple recovery sidechains competing for CRQC-recovered dormant assets. Each sidechain would accept financial risks arising from the legal obligation to return funds to the original owners and would bear the costs of acquiring dormant assets from CRQC operators while attempting to turn a profit by charging owners a deposit fee for the amount of time the recovered assets were held on the sidechain and eventually on the portion of assets for which the return obligation ultimately expires. Compared to Do Nothing, Burn, and Hourglass, the Bad Sidechain approach maximizes the value returned to the rightful owners as well as the value retained in the cryptocurrency community (assuming this community develops, deploys and operates the recovery sidechains) at the cost of greater technical, political and legal complexity.

C. Public Policy Options for the Challenge of Dormant Assets

The most important objective of the transition to quantum-secure blockchain technology is to protect ongoing onchain economic activities. It is possible, even likely, that concerted efforts will fully attain this goal before the first CRQCs come online. By contrast, the objective of preserving transaction confidentiality on privacy-preserving blockchains, such as Zcash and Monero, cannot be fully achieved due to retroactive degradation of ECDLP-protected privacy of known addresses by quantum-capable adversaries.

Another objective is to prevent rogue actors from acquiring substantial quantum-vulnerable assets. This objective is unlikely to be fully achieved by means available to the cryptocurrency communities alone. Indeed, the Bitcoin community in particular remains divided about the handling of dormant assets [26] and some options being considered,

such as Do Nothing and Hourglass discussed earlier, explicitly allow dormant assets to be acquired by actors with CRQCs. Even if the Burn policy ultimately wins on all major blockchains, its practical application may achieve only partial success due to complexities of dealing with the inevitable long tail of unmigrated assets with lost keys and in abandoned wallets. These complexities are addressed by our Bad Sidechain proposal, but its political and technical complexities might be prohibitive. Moreover, policy disagreements may potentially lead to hard forks. Some of the forked blockchains may then retain dormant assets whose fiat value, while reduced compared to the pre-fork amount, can still be substantial. Furthermore, even if disagreements and hard forks are avoided, the long time required to migrate assets to post-quantum protocols [28] may leave a window of opportunity for actors with CRQCs to acquire a portion of dormant assets.

These challenges suggest that the problem of dormant assets is unlikely to be solved completely before the first CRQCs are built. Therefore, governments may need to consider their own options for dealing with offchain consequences. Next, we explain why legally mandated “destruction” of dormant assets would be an extremely challenging and likely ineffective policy option for governments before outlining three more realistic policy responses.

1. *Regulated Destruction of Dormant Assets is Not Viable Policy*

Many legal systems around the world place constraints on the power of governments to take over or destroy private property, such as the Fourth and Fifth Amendments to the United States Constitution. However, even governments that lack or manage to overcome such constraints cannot unilaterally “destroy” dormant cryptocurrency assets. Due to the cryptographic relationships between the blocks on the chain, rewriting the contents of a block requires changing all the blocks that follow it. However, some dormant assets, in particular an overwhelming majority of the P2PK coins, are recorded in some of the oldest blocks of the chain (see Figure 4), so erasing them would entail rewriting essentially the entire blockchain history. The computational cost of doing so is equivalent to the computational cost of re-mining more than a decade of Bitcoin’s history.

However, a government could require Bitcoin miners in their jurisdiction to reject transactions that spend dormant assets. The cryptocurrency community refers to this as transaction censorship and regards it as extremely controversial. Indeed, removing the need for third party intermediaries together with the associated costs and risks, including transaction censorship, are among Bitcoin’s central design objectives [26, 140, 305]. Blockchain networks exhibit a significant degree of censorship resistance due to the decentralized nature of their architecture, transaction settlement mechanism and governance. If a government banned mining transactions spending dormant assets, then foreign miners outside cooperating jurisdictions could continue to include them in their blocks. Thus, this policy action can delay and increase costs for offending transactions, but does not prevent them from ultimately being settled on the blockchain. Indeed, this is what happened after the U.S. Treasury’s Office of Foreign Assets Control (OFAC) sanctioned Tornado Cash [70, 306], a smart contract protocol, in August 2022 [307, 308]. The sanctions were successful in substantially reducing user participation in the protocol and delayed its transactions as many block builders and validators chose to cooperate. However, they ultimately failed to prevent Tornado Cash transactions from eventually being settled by uncooperating network participants.

Therefore, a government would in fact have to require domestic miners not only to reject transactions that spend dormant assets but also valid blocks mined elsewhere and containing banned transactions. Such an action would likely result in Bitcoin mining moving abroad, as it did in 2021 following the ban on Bitcoin mining in China [309–311]. Alternatively, if a sufficiently large number of miners stayed, the blockchain would undergo a hard fork splitting into two — a domestically mined censored chain and a foreign-mined uncensored one — while leaving every owner with an equal amount of digital assets on both forks. Ultimately, the action would solve the challenge of dormant assets on the censored domestic blockchain while leaving it unsolved on the uncensored one. The degree of success would depend on the relative value of the two cryptocurrencies which would in turn depend on the behavior of market participants, especially large cryptocurrency holders. Given the importance of censorship resistance to cryptocurrency communities, it is reasonable to expect the value of the censored cryptocurrency to suffer significantly due to loss of trust and abandonment by users. In addition, offchain, the action could trigger litigation against expropriation and destruction of billions of USD of private property. At the end of the day, despite the high costs, legally-mandated transaction censorship would likely fail to secure dormant assets against quantum attacks as they would remain vulnerable on the uncensored foreign-mined blockchain.

In order for this type of forceful policy response to have any hope of success, a large fraction of all miners globally would have to acquiesce to transaction censorship. Rough estimates put the joint mining power of the United States, China, Kazakhstan, Canada and the European Union at about 85% of global mining power [312] (these estimates must be treated with caution since reliable information about miner location is not easily available). This falls short of majorities generally associated with consensus-based protocol upgrades, such as Taproot which reached 99.85% of miner support [313]. Moreover, despite the estimates that miners are concentrated in the above countries, the

cryptocurrency mining industry is mobile and responsive to financial incentives [309], so in order to deny rogue actors equipped with CRQCs access to dormant assets, a broad, multi-lateral government response would likely be needed to coordinate a long-lasting impact. However, each individual government would be incentivized to allow the proceeds from liquidation of these assets to be realized in their jurisdiction, creating the dynamics of a multi-party Prisoner’s Dilemma [314, 315]. This would constrain and complicate international cooperation [316] while traditional foreign policy tools for overcoming these limits and complications via institution building [317] would likely prove inadequate due to the possibility of state support for some CRQC-equipped rogue actors. Given the inherent challenges of tightly enforced and sustained international regulation of a globally dispersed and mobile industry, this policy action would face immense practical barriers to success, in addition to forceful opposition from the cryptocurrency community.

2. Digital Salvage

Bitcoin’s ownership model (as well as that of many other cryptocurrencies) is based on the principle that the owner of any given asset is anyone possessing the knowledge of the private key corresponding to the public key to which the asset is locked on the ledger. This principle rests on the assumption that the private key cannot be computed by anyone else. However, a CRQC can quickly derive a private key from any public key, thus breaking the foundation of Bitcoin’s model of ownership. As soon as actors with access to CRQCs enter the scene, confidence that parties transacting in Bitcoin once had in the ownership of their respective digital assets is irrevocably lost. At this point, none of Bitcoin’s onchain mechanisms provide a reliable means of identifying asset owners.

One public policy option treats dormant assets as effectively abandoned property subject to regulated digital salvage (akin to “sunken treasure”). Existing legal frameworks and property laws provide general guidance for dealing with various types of lost, abandoned and unclaimed assets and help imagine what digital salvage might look like. Principles such as escheatment [318], treasure trove, adverse possession, and maritime salvage offer general precedent for regulated recovery of abandoned property. A legal argument that dormant assets are lost or unowned assets can be made based on the fact that none of the information on the Bitcoin blockchain enables one to distinguish the true owner from someone with access to a CRQC (though in some cases offchain information may fill that gap). Another general legal principle that may apply is the doctrine of laches [319]; if an owner of dormant coins has known for years that their assets are at risk and has failed to transact them to a post-quantum address, then they may be deemed to have failed to assert their rights through inaction.

In the United States, a specific law that might provide a concrete policy blueprint is the Revised Uniform Unclaimed Property Act (2016) [320] which governs the handling of abandoned property, both tangible and intangible, including “virtual currency”. In particular, its Article 10 regulates examinations of records for the purpose of identification of unclaimed assets and the transfer of their custody to state authorities. RUUPA implementations specify inactivity time necessary for an asset to be deemed abandoned and regulate fees for record examinations. Following the logic of Article 10, a contract auditor might assist a government’s administrator in examination of transaction records on the blockchain to identify digital assets whose public keys have been exposed and which therefore lack reliably identifiable owners. Subsequently, the auditor could use a CRQC to derive the private keys of the identified assets to facilitate the transfer of their custody to state authorities. This type of policy option would introduce a powerful, market-based incentive for quantum computer development and would allow for the resulting financial gains to be transparently integrated into the formal, taxable economy, before they are recovered by bad actors.

However, even though these laws and principles provide a potential basis to design a digital salvage policy, none of them meets the challenge in their present form. In particular, RUUPA deals with assets in possession and control of a business entity called a “holder” and distinct from the assets’ owner. However, no party involved in the operation of the Bitcoin blockchain clearly meets the legal requirements to be the “holder” of the dormant coins. Indeed, none of them possess or control the coins since none of them know the private key. Moreover, RUUPA is a model law of which individual US states enact their own variants while the Bitcoin blockchain is a single world-wide distributed ledger. In the absence of legal precedent, it is ambiguous as to whether the Bitcoin blockchain falls under the purview of RUUPA and it is unclear how one would resolve inevitable conflicts in its application in different states.

3. National Security Response

The transnational character of the Bitcoin blockchain suggests a second type of policy option that recognizes the spectre of dormant assets falling to rogue actors as a national security risk. Indeed, some governments will have the option of using CRQCs (or paying a bounty to companies) to acquire these assets (possibly to burn them by sending them to the unspendable `OP_RETURN` address [321]) as a national security matter. As before, blockchain’s loss of the

ability to reliably identify asset owners combined with the laches doctrine [319] enables governments to argue that the original owners, through years of inaction, have failed to assert their property rights.

4. *Engagement with Bitcoin Community*

Finally, governments have the option to engage in a dialog with the cryptocurrency community to advocate for solutions based on their broader offchain consequences. Bitcoin has its origins in the cypherpunk movement espousing the “code is law” philosophy [322] and traditionally characterized by a techno-libertarian ethos [323]. The community fiercely guards its independence and has a strong preference for reasoned technical solutions arrived at through consensus-seeking discussion over forceful legal interventions. Its general skepticism of government intervention has also been expressed by Bitcoin’s pseudonymous creator Satoshi Nakamoto in the genesis block (block zero) which immortalizes a news headline: “The Times 03/Jan/2009 Chancellor on brink of second bailout for banks”. All of this suggests that the Bitcoin community may prove a potentially challenging partner for governments. However, in recent years, as the Bitcoin community has come to resemble the wider society, the original techno-libertarian ethos has lost some of its former centrality potentially making them easier for governments to engage with it. Moreover, the community has always had a strong sense of the ecosystem’s own interests, including its financial stability and continued success, suggesting potentially shared interests with governments, especially for large cryptocurrency holders. This creates the possibility of a productive engagement and an opportunity for governments to shape the community’s ongoing debates about the challenge of P2PK coins and other dormant assets.

IX. OUTLOOK

The emergence of CRQCs represents a serious threat to cryptocurrencies that demands a close examination of possible developments at the intersection of quantum computing and digital finance. While the quantum computing and cryptocurrency communities have largely operated in isolation, the significant reduction in resource requirements detailed here necessitates a convergence of these two worlds. Our finding that a superconducting architecture could break 256-bit ECDLP with fewer than half a million physical qubits challenges conventional wisdom regarding the timeline of the threat. This is not merely a distant danger to dormant keys; the potential for early fast-clock CRQCs to launch on-spend attacks within Bitcoin’s ten minute average block time places active transactions at immediate risk, dismantling the assumption that fee-based prioritization alone can outrun quantum adversaries.

This rapidly closing window forces the Bitcoin community to face urgent and difficult decisions regarding legacy assets, such as the 1.7 million bitcoin locked in P2PK scripts and an even greater amount of assets vulnerable due to address reuse. Because dormant assets cannot be migrated via standard software updates, they represent a fixed, multibillion-dollar target that will inevitably become accessible to a quantum actor. One way to address this would be a bifurcated approach that respects the distinct authorities of the state and the decentralized network. Governments have a range of policy options to prevent these assets from falling to rogue actors, such as regulated salvage, national security intervention, and engagement with cryptocurrency communities to shape technical solutions. These policy actions help enable the funds resulting from the recovery of P2PK coins and other dormant assets to flow into the formal, taxable economy rather than empowering political adversaries, rogue states, or criminal syndicates. A clear policy framework could act as a complement to, rather than a substitute for, technical intervention. The Bitcoin community retains the autonomy to decide whether to “burn” these vulnerable coins to preserve the chain’s financial stability, allowing these strategies to function in parallel: governments create legal containers to manage the potential chaos of inevitable salvage, while the community simultaneously weighs the difficult choice of breaking immutability to destroy the vulnerable supply.

Ethereum also faces substantial at-rest vulnerabilities that must be dealt with before CRQCs arrive. The account model uses vulnerable elliptic curves as a core component of onchain identity, putting all accounts that have carried out transactions at risk including high value accounts, such as exchange hot wallets. Smart contracts with exposed admin keys that cannot be easily rotated (without draining and replacing the contracts themselves) create a logistical bottleneck for security upgrades that puts “low ether, high leverage” accounts and contracts responsible for tokenized real-world assets, oracles, bridges, guardians, etc. at risk. Moreover, the potential compromise of validators threatens the integrity of the Proof-of-Stake consensus mechanism itself, creating an existential risk to the chain’s continued operation. Finally, the vulnerability of Data Availability Sampling mechanism opens it up to on-setup attacks that can be launched without a quantum computer using a reusable exploit created once on a CRQC. Ethereum is more exposed than Bitcoin due to the prevalence of at-rest vulnerabilities, but its recent active steps towards PQC migration promise a potentially more expedient transition to quantum-safe protocols.

Ongoing developments in other blockchain technologies and financial services also add to the urgency. In particular, tokenization of real-world assets, actively taking place on blockchains such as Ethereum and Solana, is expected to open up markets projected to exceed 16 trillion USD by 2030 [51]. At the same time, the cryptographic protocols underlying tokenization remain vulnerable to quantum attacks. In fact, upgrades and new features continue to be built on quantum-vulnerable cryptography and in some cases even weaken existing protections. Given that tokenization is projected to push the potential value of quantum-vulnerable assets above the “too-big-to-fail” economic stability thresholds, the situation may justify policy interventions, such as a mandate to deploy post-quantum cryptography.

We contend that the amount of time remaining before the arrival of CRQCs still exceeds the amount of time needed to migrate public blockchains to PQC, though the margin for error is increasingly narrow. Therefore, we have offered updated resource estimates for quantum attacks on blockchain cryptography together with an analysis of vulnerabilities and mitigations in order to urge all vulnerable cryptocurrency communities to begin PQC transition immediately while its timely completion is still the likely prospect. We have also discussed policy options to address offchain consequences of some of the possible migration paths. Finally, we have drawn attention to the numerous ongoing efforts towards PQC transition, such as the development of post-quantum blockchains like the QRL and Abelian, the integration of post-quantum protocols on Algorand, post-quantum experimentation on Solana and the XRP Ledger, as well as the active research and development initiatives spearheaded by the Ethereum Foundation. Above all, these trailblazing projects demonstrate that transition to post-quantum cryptography is realistic and instill hope that it will have been completed before the first CRQCs come online.

Acknowledgements

We thank Greg Kahanamoku-Meyer and Martin Ekerå for discussions related to the cryptanalysis resource estimates presented in this work. We thank Alan Ho, Alex Pruden, Ansgar Dietrichs, Conor Deegan, Pierre-Luc Dallaire-Demers and Liam Horne for discussions about the quantum vulnerabilities of cryptocurrencies. We thank Mark Zhandry and Alex Meiburg for discussions about our approach to zero-knowledge proofs. We thank James Manyika, Kate Weber, Sergio Boixo, Sophie Schmiege and Thomas Plumb-Reyes for careful review and feedback on this whitepaper.

Financial Conflict of Interest Statement

We the authors attest that at the time of the initial arXiv and IACR ePrint publication of this article, none of us hold any short positions against any cryptocurrency assets. Some of us hold long positions in cryptocurrencies, including some that involve the use of post-quantum cryptography. The authors reserve the right to initiate any positions in these assets in the future.

-
- [1] R. P. Feynman, Simulating physics with computers, *International Journal of Theoretical Physics* **21** (1982).
 - [2] P. Schleich, L. M. Calderón, C. Sun, M. Bagherimehrab, A. Aldossary, J. S. Kottmann, and A. Aspuru-Guzik, *Quantum Computing for Quantum Chemistry* (American Chemical Society, Washington, DC, USA, 2025).
 - [3] R. Babbush, R. King, S. Boixo, W. Huggins, T. Khattar, G. H. Low, J. R. McClean, T. O’Brien, and N. Rubin, The grand challenge of quantum applications, [arXiv:2511.09124](https://arxiv.org/abs/2511.09124) (2025).
 - [4] A. Montanaro, Quantum algorithms: an overview, *npj Quantum Information* **2**, 15023 (2016).
 - [5] J. Chou, J. Manyika, and H. Neven, The race to lead the quantum future, *Foreign Affairs* (2025).
 - [6] D. J. Bernstein and T. Lange, Post-quantum cryptography, *Nature* **549**, 188 (2017).
 - [7] D. Joseph, R. Misoczki, M. Manzano, J. Tricot, F. D. Pinuaga, O. Lacombe, S. Leichenauer, J. Hidary, P. Venables, and R. Hansen, Transitioning organizations to post-quantum cryptography, *Nature* **605**, 237 (2022).
 - [8] C. Gidney, How to factor 2048 bit RSA integers with less than a million noisy qubits, [arXiv:2505.15917](https://arxiv.org/abs/2505.15917) (2025).
 - [9] R. C. Merkle, Secure communications over insecure channels, *Commun. ACM* **21**, 294 (1978).
 - [10] R. L. Rivest, A. Shamir, and L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, *Communications of the ACM* **21**, 120 (1978).
 - [11] V. S. Miller, Use of elliptic curves in cryptography, in *Advances in Cryptology — CRYPTO ’85 Proceedings*, edited by H. C. Williams (Springer Berlin Heidelberg, Berlin, Heidelberg, 1986) pp. 417–426.
 - [12] N. Koblitz, Elliptic curve cryptosystems, *Mathematics of Computation* **48**, 203 (1987).
 - [13] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM Journal on Computing* **26**, 1484 (1997).
 - [14] AMD, Google, Meta, Microsoft, and NVIDIA, *OCP GPU FW Update Specification*, Tech. Rep. (Open Compute Project, 2023).

- [15] K. Schneider, L. Auer, and A. Wagner, Fault attacks on ECC signature verification, *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2025**, 1010 (2025).
- [16] E. Rescorla, *The Transport Layer Security (TLS) Protocol Version 1.3*, RFC 8446 (2018).
- [17] M. Marlinspike and T. Perrin, *The X3DH key agreement protocol* (2016).
- [18] J. W. Bos, J. A. Halderman, N. Heninger, J. Moore, M. Naehrig, and E. Wustrow, *Elliptic curve cryptography in practice*, Cryptology ePrint Archive, Paper 2013/734 (2013).
- [19] *YubiKey Technical Manual*, Yubico (2026).
- [20] D. Stebila and J. Green, *Elliptic Curve Algorithm Integration in the Secure Shell (SSH) Transport Layer*, RFC 5656 (2009).
- [21] P. Hoffman and W. Wijngaards, *Elliptic Curve Digital Signature Algorithm (DSA) for DNSSEC*, RFC 6605 (2012).
- [22] N. Gura, A. Patel, A. Wander, H. Eberle, and S. C. Shantz, Comparing elliptic curve cryptography and RSA on 8-bit CPUs, in *Cryptographic Hardware and Embedded Systems - CHES 2004*, edited by M. Joye and J.-J. Quisquater (Springer Berlin Heidelberg, Berlin, Heidelberg, 2004) pp. 119–132.
- [23] C. A. Lara-Nino, A. Diaz-Perez, and M. Morales-Sandoval, Elliptic curve lightweight cryptography: A survey, *IEEE Access* **6**, 72514 (2018).
- [24] E. Barker, *NIST SP 800-57 Part 1 Rev. 5: Recommendation for Key Management: Part 1 – General*, NIST Special Publication 800-57 Part 1 Rev. 5 (National Institute of Standards and Technology, Gaithersburg, MD, USA, 2020).
- [25] A. Milton and C. Shikhelman, *Bitcoin and Quantum Computing: Current Status and Future Directions*, Research Report (Chaincode Labs, New York, NY, 2025) accessed: 2026-03-22.
- [26] A. Li *et al.*, *The Quantum Threat to Bitcoin*, Human Rights Research Report (Human Rights Foundation, 2025) informed by technical insights from the 2025 Presidio Bitcoin Quantum Summit.
- [27] C. Deegan, *Quantum vulnerability of bitcoin addresses*, Project Eleven Blog (2025).
- [28] J. J. Pont, J. J. Kearney, J. Moyler, and C. A. Perez-Delgado, Downtime required for bitcoin quantum-safety, [arXiv:2410.16965](https://arxiv.org/abs/2410.16965) (2024).
- [29] D. Aggarwal, G. Brennen, T. Lee, M. Santha, and M. Tomamichel, Quantum attacks on bitcoin, and how to protect against them, *Ledger* **3**, 10.5195/ledger.2018.127 (2018).
- [30] T. Ruffing, *The post-quantum security of bitcoin’s taproot as a commitment scheme*, Cryptology ePrint Archive, Paper 2025/1307 (2025).
- [31] S. Holmes and L. Chen, *Assessment of quantum threat to bitcoin and derived cryptocurrencies*, Cryptology ePrint Archive, Paper 2021/967 (2021).
- [32] K. Fukuda, S. Matsuo, Y. Suga, and T. Ito, *The grand challenge of PQC migration: Analysis of modern blockchain and intertwined human egoisms*, Cryptology ePrint Archive, Paper 2025/1626 (2025).
- [33] D. B. C. Costa, *Post-Quantum Financial Infrastructure Framework (PQFIF): A Roadmap for the Quantum-Safe Transition of Global Financial Infrastructure*, Regulatory Submission (U.S. Securities and Exchange Commission (SEC), 2025) prepared for the U.S. Crypto Assets Task Force – SEC.
- [34] A. Maurushat, *Disclosure of Security Vulnerabilities: Legal and Ethical Issues* (Springer London, 2013).
- [35] ISO/IEC, *ISO/IEC 29147:2018 Information technology — Security techniques — Vulnerability disclosure*, Standard (International Organization for Standardization and International Electrotechnical Commission, Geneva, Switzerland, 2018).
- [36] D. C. North, Institutions, *Journal of Economic Perspectives* **5**, 97–112 (1991).
- [37] D. C. North, *Institutions, Institutional Change and Economic Performance*, Political Economy of Institutions and Decisions (Cambridge University Press, 1990).
- [38] CoinMarketCap, *FUD*, CoinMarketCap Glossary (2021), accessed 2026-03-28.
- [39] BTSE, *Crypto Trading Psychology: Dealing with FUD and FOMO in the Cryptocurrency Market*, Special Report (BTSE, 2023) BTSE Trading Psychology Series. Accessed March 28, 2026.
- [40] R. Harris and A. McIntyre, *The Complete Sales Letter Book: Model Letters for Every Selling Situation*, Sharpe Professional (Sharpe Professional, 1998).
- [41] S. Goldwasser, S. Micali, and C. Rackoff, The knowledge complexity of interactive proof-systems, in *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC ’85 (Association for Computing Machinery, New York, NY, USA, 1985) pp. 291–304.
- [42] J.-J. Quisquater, M. Quisquater, M. Quisquater, M. Quisquater, L. C. Guillou, M. A. Guillou, G. Guillou, A.-C. Guillou, G. Guillou, S. Guillou, and T. A. Berson, How to explain zero-knowledge protocols to your children, in *Annual International Cryptology Conference* (1989).
- [43] S. Cu’ellar Gempeler, B. Harris, J. Parker, S. Pernsteiner, I. Sweet, and E. Tromer, Cheesecloth: Zero-knowledge proofs of real-world vulnerabilities, *ACM Trans. Priv. Secur.* **28**, 10.1145/3747589 (2025).
- [44] D. Litinski, How to compute a 256-bit elliptic curve private key with only 50 million toffoli gates, [arXiv:2306.08585](https://arxiv.org/abs/2306.08585) (2023).
- [45] C. Gidney and M. Ekerå, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, *Quantum* **5**, 433 (2021).
- [46] C. Gidney, How to factor 2048 bit RSA integers with less than a million noisy qubits, [arXiv:2505.15917](https://arxiv.org/abs/2505.15917) (2025).
- [47] M. Habov stiak, *Hashed keys are actually fully quantum secure*, Bitcoin Development Mailing List (Google Groups) (2025), mailing list thread featuring contributions from Lloyd Fournier, Agustin Cruz, Antoine Poinso, and others. Proposes a commit-reveal scheme to protect hashed pubkeys from mempool-snatching by quantum adversaries.
- [48] D. J. Bernstein, Introduction to post-quantum cryptography, in *Post-Quantum Cryptography*, edited by D. J. Bernstein, J. Buchmann, and E. Dahmen (Springer Berlin Heidelberg, Berlin, Heidelberg, 2009) pp. 1–14.

- [49] D. Bernstein and T. Lange, *Post-quantum cryptography: dealing with the fallout of physics success*, Cryptology ePrint Archive (International Association for Cryptologic Research, 2017).
- [50] National Institute of Standards and Technology, Glossary: Defense-in-depth, https://csrc.nist.gov/glossary/term/defense_in_depth (2026), CSRC Computer Security Resource Center.
- [51] P. Gretz, D. Gross, E. M. Haynes, Y. Tarek, A. Toma, A. Chitgupi, and T. Kurowski, *Tokenization Standards: Taming the Regulatory Menagerie*, Industry Research Report (Nethermind and PwC Germany, 2025) distributed in collaboration with Global Finance & Technology Network (GFTN).
- [52] A. Yakovenko, *Solana: A New Architecture for a High Performance Blockchain*, Protocol White Paper v0.8.13 (Solana Labs, 2017).
- [53] J. Chen and S. Micali, Algorand, [arXiv:1607.01341](https://arxiv.org/abs/1607.01341) (2016).
- [54] M. Likhava, G. Losa, D. Mazières, G. Hoare, N. Barry, E. Gafni, J. Jove, R. Malinowsky, and J. McCaleb, Fast and secure global payments with stellar, in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, SOSP '19 (Association for Computing Machinery, New York, NY, USA, 2019) p. 80–96.
- [55] D. Mazières, *The Stellar Consensus Protocol: A Federated Model for Internet-level Consensus*, Tech. Rep. (Stellar Development Foundation, 2015).
- [56] B. Chase and E. MacBrough, Analysis of the xrp ledger consensus protocol, [arXiv:1802.07242](https://arxiv.org/abs/1802.07242) (2018).
- [57] Ripple Developers, *rippled: XRP ledger network server*, GitHub Repository (2026).
- [58] Algorand Foundation, *Asset tokenization on Algorand: Powering a more open financial system*, Official Ecosystem Portal (2026).
- [59] Stellar Development Foundation, *RWA tokenization*, Official Ecosystem Portal (2026).
- [60] Team RippleX, *The Future of Asset Tokenization: A New Token Standard for Institutional-Grade Finance on XRP Ledger*, Industry Research Report (Ripple Labs, Inc., 2025).
- [61] D. O'Brien, *Protecting Chrome traffic with Hybrid Kyber KEM*, Chromium Blog (2023), accessed March 26, 2026.
- [62] L. Young, *Technical brief: Quantum-resistant transactions on Algorand with Falcon signatures*, Algorand Developer Blog (2025), technical contributions by Giulio Pizzini, Cosimo Bassi, and Steve Ferrigno.
- [63] The QRL Foundation, *Quantum Resistant Ledger: A visionary, future-proof blockchain with unparalleled security* (2026).
- [64] P. Waterland, *Quantum Resistant Ledger (QRL)*, Technical White Paper (The QRL Foundation, 2016).
- [65] M. Zweil, *Mochimo: Post-Quantum Currency*, Protocol White Paper (Adequate Systems, LLC, 2018).
- [66] Alice, Bob, Eve, and Lambda, *Abelian (ABEL) – A Quantum-Resistant Cryptocurrency Balancing Privacy and Accountability*, Cryptographic White Paper (The Abelian Foundation, 2022).
- [67] ForkLog, *XRP Ledger implements quantum threat protection*, ForkLog Magazine (2025).
- [68] C. Deegan, *A look at post quantum proposals for bitcoin*, Project Eleven Blog (2025).
- [69] F. Wang, S. Cohny, and J. Bonneau, *SoK: Trusted setups for powers-of-tau strings*, Cryptology ePrint Archive, Paper 2025/064 (2025).
- [70] A. Pertsev, R. Semenov, and R. Storm, *Tornado Cash Privacy Solution*, Technical Whitepaper Version 1.4 (Tornado Cash, 2019) accessed via Berkeley DeFi educational resources.
- [71] R. Acharya, D. A. Abanin, L. Aghababaie-Beni, I. Aleiner, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, N. Astrakhantsev, J. Atalaya, R. Babbush, D. Bacon, B. Ballard, J. C. Bardin, J. Bausch, A. Bengtsson, A. Bilmes, S. Blackwell, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, D. A. Browne, B. Buchea, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, A. Cabrera, J. Campero, H.-S. Chang, Y. Chen, Z. Chen, B. Chiaro, D. Chik, C. Chou, J. Claes, A. Y. Cleland, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, S. Das, A. Davies, L. De Lorenzo, D. M. Debroy, S. Demura, M. Devoret, A. Di Paolo, P. Donohoe, I. Drozdov, A. Dunsworth, C. Earle, T. Edlich, A. Eickbusch, A. M. Elbag, M. Elzouka, C. Erickson, L. Faoro, E. Farhi, V. S. Ferreira, L. F. Burgos, E. Forati, A. G. Fowler, B. Foxen, S. Ganjam, G. Garcia, R. Gasca, É. Genois, W. Giang, C. Gidney, D. Gilboa, R. Gosula, A. G. Dau, D. Graumann, A. Greene, J. A. Gross, S. Habegger, J. Hall, M. C. Hamilton, M. Hansen, M. P. Harrigan, S. D. Harrington, F. J. H. Heras, S. Heslin, P. Heu, O. Higgott, G. Hill, J. Hilton, G. Holland, S. Hong, H.-Y. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, S. Jordan, C. Joshi, P. Juhas, D. Kafri, H. Kang, A. H. Karamlou, K. Kechedzhi, J. Kelly, T. Khaire, T. Khattar, M. Khezri, S. Kim, P. V. Klimov, A. R. Klots, B. Kobrin, P. Kohli, A. N. Korotkov, F. Kostritsa, R. Kothari, B. Kozlovskii, J. M. Kreikebaum, V. D. Kurilovich, N. Lacroix, D. Landhuis, T. Lange-Dei, B. W. Langley, P. Laptev, K.-M. Lau, L. Le Guevel, J. Ledford, J. Lee, K. Lee, Y. D. Lensky, S. Leon, B. J. Lester, W. Y. Li, Y. Li, A. T. Lill, W. Liu, W. P. Livingston, A. Locharla, E. Lucero, D. Lundahl, A. Lunt, S. Madhuk, F. D. Malone, A. Maloney, S. Mandrá, J. Manyika, L. S. Martin, O. Martin, S. Martin, C. Maxfield, J. R. McClean, M. McEwen, S. Meeks, A. Megrant, X. Mi, K. C. Miao, A. Mieszala, R. Molavi, S. Molina, S. Montazeri, A. Morvan, R. Movassagh, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, C.-H. Ni, M. Y. Niu, T. E. O'Brien, W. D. Oliver, A. Opremcak, K. Ottosson, A. Petukhov, A. Pizzuto, J. Platt, R. Potter, O. Pritchard, L. P. Pryadko, C. Quintana, G. Ramachandran, M. J. Reager, J. Redding, D. M. Rhodes, G. Roberts, E. Rosenberg, E. Rosenfeld, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, A. W. Senior, M. J. Shearn, A. Shorter, N. Shutty, V. Shvarts, S. Singh, V. Sivak, J. Skrzynny, S. Small, V. Smelyanskiy, W. C. Smith, R. D. Somma, S. Springer, G. Sterling, D. Strain, J. Suchard, A. Szasz, A. Szein, D. Thor, A. Torres, M. M. Torunbalci, A. Vaishnav, J. Vargas, S. Vdovichev, G. Vidal, B. Villalonga, C. V. Heidweiller, S. Waltman, S. X. Wang, B. Ware, K. Weber, T. Weidel, T. White, K. Wong, B. W. K. Woo, C. Xing, Z. J. Yao, P. Yeh, B. Ying, J. Yoo, N. Yosri, G. Young, A. Zalcman, Y. Zhang, N. Zhu, and N. Zobrist, Quantum error correction below the surface code threshold, *Nature* **638**, 920 (2024).

- [72] M. L. Chan, A. A. Capatos, P. Lodahl, A. S. Sørensen, and S. Paesani, Practical blueprint for low-depth photonic quantum computing with quantum dots, [arXiv:2507.16152](https://arxiv.org/abs/2507.16152) (2025).
- [73] J. R. Scott and K. C. Balram, Timing constraints due to real-time graph-traversal algorithms on incomplete cluster states in photonic measurement-based quantum computing, *Physical Review Applied* **20**, [10.1103/physrevapplied.20.024019](https://doi.org/10.1103/physrevapplied.20.024019) (2023).
- [74] P. Stano and D. Loss, Review of performance metrics of spin qubits in gated semiconducting nanostructures, *Nature Reviews Physics* **4**, 672 (2022).
- [75] D. Bluvstein, A. A. Geim, S. H. Li, S. J. Evered, J. P. Bonilla Ataides, G. Baranes, A. Gu, T. Manovitz, M. Xu, M. Kalinowski, S. Majidy, C. Kokail, N. Maskara, E. C. Trapp, L. M. Stewart, S. Hollerith, H. Zhou, M. J. Gullans, S. F. Yelin, M. Greiner, V. Vuletić, M. Cain, and M. D. Lukin, A fault-tolerant neutral-atom architecture for universal quantum computation, *Nature* **649**, 39 (2025).
- [76] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, Fast quantum logic gates with trapped-ion qubits, *Nature* **555**, 75 (2018).
- [77] S. Aaronson, More on whether useful quantum computing is “imminent” — Shtetl-Optimized, <https://scottaaronson.blog/?p=9425> (2025), accessed: 2026-03-22.
- [78] S. Aaronson, On reducing the cost of breaking RSA-2048 to 100,000 physical qubits — Shtetl-Optimized, <https://scottaaronson.blog/?p=9425> (2026), accessed: 2026-03-22.
- [79] A. Kerckhoffs, La Cryptographie Militaire, *Journal des Sciences Militaires* **IX**, 5 (1883).
- [80] CERT Coordination Center, [CERT/CC vulnerability disclosure policy](#), Official Policy Document (2026).
- [81] Google Project Zero, [Vulnerability disclosure policy](#), Official Policy Document (2026).
- [82] N. Carter, Bitcoin developers are mostly not concerned about quantum risk – Murmurations II, <https://murmurationstwo.substack.com/p/bitcoin-developers-are-mostly-not> (2026), accessed: 2026-03-22.
- [83] Succinct Labs, [SP1: The zero-knowledge virtual machine \(zkVM\)](#), GitHub Repository (2026).
- [84] J. Jancar and V. Sedlacek, Parameters for secp256k1 — Standard Curves Database, <https://std.neuromancer.sk/secg/secp256k1> (2026), technical specifications for the Koblitz curve used by Bitcoin and Ethereum. Accessed: 2026-03-21.
- [85] D. R. L. Brown, *SEC 2: Recommended Elliptic Curve Domain Parameters*, Technical Specification Version 2.0 (Certicom Research, 2010) standards for Efficient Cryptography (SEC).
- [86] J. Proos and C. Zalka, [Shor’s discrete logarithm quantum algorithm for elliptic curves](#) (2003).
- [87] C. Cheviguard, P.-A. Fouque, and A. Schrottenloher, [Reducing the number of qubits in quantum discrete logarithms on elliptic curves](#), Cryptology ePrint Archive, Paper 2026/280 (2026).
- [88] T. Häner, S. Jaques, M. Naehrig, M. Roetteler, and M. Soeken, [Improved quantum circuits for elliptic curve discrete logarithms](#), Cryptology ePrint Archive, Paper 2020/077 (2020).
- [89] É. Gouzien, D. Ruiz, F.-M. Le Régent, J. Guillaud, and N. Sangouard, Performance analysis of a repetition cat code architecture: Computing 256-bit elliptic curve logarithm in 9 hours with 126,133 cat qubits, *Physical Review Letters* **131**, [10.1103/physrevlett.131.040602](https://doi.org/10.1103/physrevlett.131.040602) (2023).
- [90] P.-L. Dallaire-Demers, W. Doyle, and T. Foo, Brace for impact: Ecdlp challenges for quantum cryptanalysis, [arXiv:2508.14011](https://arxiv.org/abs/2508.14011) (2025).
- [91] A. Fiat and A. Shamir, How to prove yourself: Practical solutions to identification and signature problems, in *Advances in Cryptology — CRYPTO’ 86*, edited by A. M. Odlyzko (Springer Berlin Heidelberg, Berlin, Heidelberg, 1987) pp. 186–194.
- [92] R. Babbush, A. Zalcman, C. Gidney, M. Broughton, T. Khattar, H. Neven, T. Bergamaschi, J. Drake, and D. Boneh, [Data for “securing elliptic curve cryptocurrencies against quantum vulnerabilities: Resource estimates and mitigations”](#) (2026).
- [93] C. Jones, Low-overhead constructions for the fault-tolerant toffoli gate, *Physical Review A* **87**, [10.1103/physreva.87.022328](https://doi.org/10.1103/physreva.87.022328) (2013).
- [94] A. G. Fowler, Time-optimal quantum computation, [arXiv:1210.4626](https://arxiv.org/abs/1210.4626) (2012).
- [95] P. L. Montgomery, Speeding the Pollard and elliptic curve methods of factorization, *Mathematics of Computation* **95**, 243 (1987).
- [96] H. Kim, K. Jang, S. Wang, A. Baksi, G. Song, H. Seo, and A. Chattopadhyay, [New quantum circuits for ECDLP: Breaking prime elliptic curve cryptography in minutes](#), Cryptology ePrint Archive, Paper 2026/106 (2026).
- [97] M. Roetteler, M. Naehrig, K. M. Svore, and K. Lauter, Quantum resource estimates for computing elliptic curve discrete logarithms, [arXiv:1706.06752](https://arxiv.org/abs/1706.06752) (2017).
- [98] C. Gidney, M. Newman, P. Brooks, and C. Jones, Yoked surface codes, [arXiv:2312.04522](https://arxiv.org/abs/2312.04522) (2023).
- [99] C. Gidney and M. Newman, [Yoked surface codes](#), Talk at the 7th International Conference on Quantum Error Correction (QEC25), New Haven, Connecticut, USA (2025).
- [100] C. Gidney and A. G. Fowler, Flexible layout of surface code computations using autocccz states, [arXiv:1905.08916](https://arxiv.org/abs/1905.08916) (2019).
- [101] M. Ekerå, *On factoring integers, and computing discrete logarithms and orders, quantumly*, Ph.D. thesis, KTH Royal Institute of Technology, Stockholm, Sweden (2024).
- [102] C. Gidney, N. Shutty, and C. Jones, Magic state cultivation: growing t states as cheap as cnot gates, [arXiv:2409.17595](https://arxiv.org/abs/2409.17595) (2024).
- [103] R. Babbush, R. King, S. Boixo, W. Huggins, T. Khattar, G. H. Low, J. R. McClean, T. O’Brien, and N. C. Rubin, The grand challenge of quantum applications, [arXiv:2511.09124](https://arxiv.org/abs/2511.09124) (2025).
- [104] B. Schneier, New Attack on AES, https://www.schneier.com/blog/archives/2011/08/new_attack_on_a_1.html (2011), accessed: 2026-03-22.

- [105] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Elucidating reaction mechanisms on quantum computers, *Proceedings of the national academy of sciences* **114**, 7555 (2017).
- [106] Z. Li, J. Li, N. S. Dattani, C. Umrigar, and G. K. Chan, The electronic complexity of the ground-state of the femo cofactor of nitrogenase as relevant to quantum simulations, *The Journal of chemical physics* **150** (2019).
- [107] N. C. Jones, R. Van Meter, A. G. Fowler, P. L. McMahon, J. Kim, T. D. Ladd, and Y. Yamamoto, Layered architecture for quantum computing, *Physical Review X* **2**, 031007 (2012).
- [108] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Physical Review A-Atomic, Molecular, and Optical Physics* **86**, 032324 (2012).
- [109] J. O’Gorman and E. T. Campbell, Quantum computation with realistic magic-state factories, *Physical Review A* **95**, 032338 (2017).
- [110] V. Gheorghiu and M. Mosca, Benchmarking the quantum cryptanalysis of symmetric, public-key and hash-based cryptographic schemes, [arXiv:1902.02332](https://arxiv.org/abs/1902.02332) (2019).
- [111] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, *PRX quantum* **2**, 030305 (2021).
- [112] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, Quantum computing enhanced computational catalysis, *Physical Review Research* **3**, 033055 (2021).
- [113] D. Rocca, C. L. Cortes, J. F. Gonthier, P. J. Ollitrault, R. M. Parrish, G.-L. Anselmetti, M. Degroote, N. Moll, R. Santagati, and M. Streif, Reducing the runtime of fault-tolerant quantum simulations in chemistry through symmetry-compressed double factorization, *Journal of Chemical Theory and Computation* **20**, 4639 (2024).
- [114] G. H. Low, R. King, D. W. Berry, Q. Han, A. E. DePrince III, A. F. White, R. Babbush, R. D. Somma, and N. C. Rubin, Fast quantum simulation of electronic structure by spectral amplification, *Physical Review X* **15**, 041016 (2025).
- [115] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, *Nature* **627**, 778 (2024).
- [116] Q. Xu, J. P. Bonilla Ataides, C. A. Pattison, N. Raveendran, D. Bluvstein, J. Wurtz, B. Vasić, M. D. Lukin, L. Jiang, and H. Zhou, Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays, *Nature Physics* **20**, 1084 (2024).
- [117] T. J. Yoder, E. Schoute, P. Rall, E. Pritchett, J. M. Gambetta, A. W. Cross, M. Carroll, and M. E. Beverland, Tour de gross: A modular quantum computer based on bivariate bicycle codes, [arXiv:2506.03094](https://arxiv.org/abs/2506.03094) (2025).
- [118] P. Webster, L. Berent, O. Chandra, E. T. Hockings, N. Baspin, F. Thomsen, S. C. Smith, and L. Z. Cohen, The pinnacle architecture: Reducing the cost of breaking rsa-2048 to 100 000 physical qubits using quantum ldpc codes, [arXiv:2602.11457](https://arxiv.org/abs/2602.11457) (2026).
- [119] J. Olle, R. Zen, M. Puviani, and F. Marquardt, Simultaneous discovery of quantum error correction codes and encoders with a noise-aware reinforcement learning agent, *npj Quantum Information* **10**, 126 (2024).
- [120] A. Y. He and Z.-W. Liu, Discovering highly efficient low-weight quantum error-correcting codes with reinforcement learning, [arXiv:2502.14372](https://arxiv.org/abs/2502.14372) (2025).
- [121] S. Yang and R. Calderbank, Spatially-coupled qldpc codes, *Quantum* **9**, 1693 (2025).
- [122] E. Heilman, *Changes to BIP-360 — Pay to Quantum Resistant Hash (P2QRH)*, Delving Bitcoin (2025), post #29 identifies the two cases arising from consideration of CRQC clock speed.
- [123] P. Anderson and M. L. Tushman, Technological discontinuities and dominant designs: A cyclical model of technological change, *Administrative Science Quarterly* **35**, 604 (1990).
- [124] G. E. Moore, Cramming More Components onto Integrated Circuits, *Electronics* **38**, 114 (1965), reprinted in: IEEE Solid-State Circuits Newsletter, Vol. 11, No. 3, Sept. 2006.
- [125] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, D. Bacon, J. C. Bardin, J. Basso, A. Bengtsson, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. Del Toro Barba, S. Demura, A. Dunsworth, D. Eppens, C. Erickson, L. Faoro, E. Farhi, R. Fatemi, L. Flores Burgos, E. Forati, A. G. Fowler, B. Foxen, W. Giang, C. Gidney, D. Gilboa, M. Giustina, A. Grajales Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. P. Harrigan, S. D. Harrington, O. Higgott, J. Hilton, M. Hoffmann, S. Hong, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveldand, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, K. Kechedzhi, J. Kelly, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, P. V. Klimov, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. Lee, B. J. Lester, A. Lill, W. Liu, A. Locharla, E. Lucero, F. D. Malone, J. Marshall, O. Martin, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, B. Meurer Costa, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, A. Morvan, E. Mount, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O’Brien, A. Opremcak, J. Platt, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, V. Shvarts, J. Skrzuzny, V. C. Smith, G. Sterling, D. Strain, M. Szalay, A. Torres, G. Vidal, B. Villalonga, C. Vollgraff Heidweiller, T. White, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, and N. Zhu, Suppressing quantum errors by scaling a surface code logical qubit, *Nature* **614**, 676 (2023).
- [126] D. Awschalom, K. K. Berggren, H. Bernien, S. Bhave, L. D. Carr, P. Davids, S. E. Economou, D. Englund, A. Faraon, M. Fejer, S. Guha, M. V. Gustafsson, E. Hu, L. Jiang, J. Kim, B. Korzh, P. Kumar, P. G. Kwiat, M. Lončar, M. D. Lukin, D. A. Miller, C. Monroe, S. W. Nam, P. Narang, J. S. Orcutt, M. G. Raymer, A. H. Safavi-Naeini, M. Spiropulu, K. Srinivasan, S. Sun, J. Vučković, E. Waks, R. Walsworth, A. M. Weiner, and Z. Zhang, Development of quantum interconnects (quics) for next-generation information technologies, *PRX Quantum* **2**, 10.1103/prxquantum.2.017002 (2021).

- [127] P. Escofet, S. B. Rached, S. Rodrigo, C. G. Almudever, E. Alarc'ón, and S. Abadal, Interconnect fabrics for multi-core quantum processors: A context analysis, in *Proceedings of the 16th International Workshop on Network on Chip Architectures*, NoC Arc '23 (Association for Computing Machinery, New York, NY, USA, 2023) pp. 34–39.
- [128] C. A. Pattison, G. Baranes, J. P. B. Ataide, M. D. Lukin, and H. Zhou, Fast quantum interconnects via constant-rate entanglement distillation, [arXiv:2408.15936](https://arxiv.org/abs/2408.15936) (2024).
- [129] H. Beast, E. Heilman, and I. F. Duke, *BIP 360: Pay-to-Merkle-Root (P2MR)*, Bitcoin Improvement Proposal 360 (Bitcoin Core, 2024) source: <https://github.com/bitcoin/bips/blob/master/bip-0360.mediawiki>. Accessed: 2026-03-22.
- [130] C. Hoskinson, *Why We Are Building Cardano: A Subjective Approach*, Foundational Whitepaper (IOHK, 2017) accessed March 26, 2026.
- [131] A. Antonopoulos, *Mastering Bitcoin: Unlocking Digital Cryptocurrencies* (O'Reilly Media, 2014).
- [132] G. Walker, Learn Me A Bitcoin, <https://learnmeabitcoin.com> (2026), accessed: 2026-03-22.
- [133] G. Walker, Script: A mini programming language — Learn Me A Bitcoin, <https://learnmeabitcoin.com/technical/upgrades/taproot/> (2025), accessed: 2026-03-21.
- [134] P. Todd, REWARD offered for hash collisions for SHA1, SHA256, RIPEMD160 and other, Bitcointalk Forum (2013).
- [135] Bitcoin Core Developers, Segregated witness wallet development guide, Official Bitcoin Core Documentation (2026).
- [136] R. Linus, BitVM: Compute Anything on Bitcoin, <https://bitvm.org/bitvm.pdf> (2023), technical White Paper.
- [137] M. Friedenbach, Why i'm against Taproot and recommend against activation on Bitcoin, Freicoin Substack (2021).
- [138] P. Wuille, J. Nick, and A. Towns, BIP 342: Validation of Taproot scripts, Bitcoin Improvement Proposals (2020), status: Deployed. Consensus Layer (Soft Fork). Requires BIP 340 and BIP 341. Licensed under BSD-3-Clause.
- [139] A. Day and C. Bookman, *Bitcoin in BigQuery: blockchain analytics on public data*, Google Cloud Blog (2018), developed in collaboration with Kaggle.
- [140] S. Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System, <https://bitcoin.org/bitcoin.pdf> (2008).
- [141] CoinPhoton, *Bitcoin network clears congestion after bullish rally*, Binance Square (2024), documents network congestion between December 2023 and January 2024 with a peak of over 194,000 unconfirmed transactions.
- [142] G. Walker, Block Reward: Incentive for miners to mine blocks — Learn Me A Bitcoin, <https://learnmeabitcoin.com/technical/upgrades/taproot/> (2025), accessed: 2026-03-21.
- [143] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability, in *2020 IEEE Symposium on Security and Privacy (SP)* (2020) pp. 910–927.
- [144] A. Chon'e and M. D. I. Canals, *Maximal Extractable Value: Implications for crypto markets*, Regulatory Risk Analysis ESMA50-481369926-29744 (European Securities and Markets Authority, 2025).
- [145] S. Nichols, *Marathon launches Slipstream tech stack to process non-standard Bitcoin transactions*, Bitcoin Magazine (via Nasdaq) (2024).
- [146] Project 11, *Bitcoin Risq List*, Technical Database (2026), a comprehensive audit of Bitcoin addresses with exposed public keys.
- [147] R. Stütz, N. Stifter, M. Dragaschnig, B. Haslhofer, and A. Judmayer, Reuse of public keys across utxo and account-based cryptocurrencies, [arXiv:2601.19500](https://arxiv.org/abs/2601.19500) (2026).
- [148] P. Wuille, *BIP 32: Hierarchical Deterministic Wallets*, Bitcoin Improvement Proposal 32 (Bitcoin Core, 2012) source: <https://github.com/bitcoin/bips/blob/master/bip-0032.mediawiki>. Accessed: 2026-03-22.
- [149] O. Sattath, On the insecurity of quantum bitcoin mining, *International Journal of Information Security* **19**, 291 (2020).
- [150] R. R. Nerem and D. R. Gaur, Conditions for advantageous quantum bitcoin mining, *Blockchain: Research and Applications* **4**, 100141 (2023).
- [151] N. Vescovo, *The coming quantum threat to Bitcoin & Rootstock: Risks, readiness, and roadmap*, Rootstock Labs Blog (2025).
- [152] E. Musk, *@grok estimate the probability of quantum computing cracking SHA-256*, X (formerly Twitter) (2025).
- [153] S. Neagle, How should Bitcoiners view quantum computing?, *Bitcoin Magazine* (2025), technical Analysis. Accessed March 29, 2026.
- [154] R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven, Focus beyond quadratic speedups for error-corrected quantum advantage, *PRX Quantum* **2**, 10.1103/prxquantum.2.010103 (2021).
- [155] C. Zalka, Grover's quantum searching algorithm is optimal, *Physical Review A* **60**, 2746 (1999).
- [156] Bitmain, *Antminer S19 Pro Specifications*, Bitmain Technologies Ltd. (2021), created Feb 27, 2020. Updated May 27, 2021.
- [157] M. Amy, O. D. Matteo, V. Gheorghiu, M. Mosca, A. Parent, and J. Schanck, *Estimating the cost of generic quantum pre-image attacks on SHA-2 and SHA-3*, Cryptology ePrint Archive, Paper 2016/992 (2016).
- [158] D. I. Ilie, S. M. Werner, I. Stewart, and W. J. Knottenbelt, Unstable throughput: When the difficulty algorithm breaks, [arXiv:2006.03044](https://arxiv.org/abs/2006.03044) (2020).
- [159] D. Boneh, B. Lynn, and H. Shacham, Short signatures from the weil pairing, in *Advances in Cryptology – ASIACRYPT 2001* (Springer Berlin Heidelberg, 2001) pp. 514–532.
- [160] C. Reitwiessner, *EIP-196: Precompiled contracts for addition and scalar multiplication on the elliptic curve altbn128*, Ethereum Improvement Proposal 196 (Ethereum Foundation, 2017) accessed: 2026-03-23.
- [161] P. S. L. M. Barreto and M. Naehrig, *Pairing-friendly elliptic curves of prime order*, Cryptology ePrint Archive, Paper 2005/133 (2005).
- [162] J. Jancar and V. Sedlacek, Parameters for BLS12-381 — Standard Curves Database, <https://std.neuromancer.sk/bls/BLS12-381> (2026), technical specifications for the BLS12-381 pairing-friendly elliptic curve. Accessed: 2026-03-21.

- [163] P. S. L. M. Barreto, B. Lynn, and M. Scott, [Constructing elliptic curves with prescribed embedding degrees](#), Cryptology ePrint Archive, Paper 2002/088 (2002).
- [164] W. Diffie and M. Hellman, New directions in cryptography, [IEEE Transactions on Information Theory](#) **22**, 644 (1976).
- [165] D.-E. Hopwood, S. Bowe, T. Hornby, and N. Wilcox, [Zcash Protocol Specification, Version 2025.6.3-84-gf58f5c \[NU6\]](#), Protocol Specification (Electric Coin Co., 2026) accessed: 2026-03-21.
- [166] koe, K. M. Alonso, and S. Noether, [Zero to Monero: Second Edition — a technical guide to a private digital currency; for beginners, amateurs, and experts](#), 2nd ed., v2.0.0 (Monero Research Lab, 2020).
- [167] D. Burkett, [LIP 004: One-Sided Transactions in MimbleWimble \(Consensus layer\)](#), Litecoin Improvement Proposal 4 (Litecoin Project, 2020) accessed: 2026-03-22.
- [168] T. Elgamal, A public key cryptosystem and a signature scheme based on discrete logarithms, [IEEE Transactions on Information Theory](#) **31**, 469 (1985).
- [169] T. P. Pedersen, Non-interactive and information-theoretic secure verifiable secret sharing, in [Advances in Cryptology — CRYPTO '91](#), edited by J. Feigenbaum (Springer Berlin Heidelberg, Berlin, Heidelberg, 1992) pp. 129–140.
- [170] T. Ruffing and G. Malavolta, [Switch commitments: A safety switch for confidential transactions](#), Cryptology ePrint Archive, Paper 2017/237 (2017).
- [171] D. Pointcheval, [Commitment schemes \(TD 16\)](#), Course Material, Master Parisien de Recherche en Informatique (MPRI) (2017), tutorial covering formal properties of Pedersen and ElGamal commitments.
- [172] A. Kate, G. M. Zaverucha, and I. Goldberg, Constant-size commitments to polynomials and their applications, in [Advances in Cryptology - ASIACRYPT 2010](#), edited by M. Abe (Springer Berlin Heidelberg, Berlin, Heidelberg, 2010) pp. 177–194.
- [173] D. Ryan, D. Feist, F. D’Amato, H.-W. Wang, and A. Stokes, [EIP-7594: PeerDAS - Peer Data Availability Sampling](#), Ethereum Improvement Proposal 7594 (Ethereum Foundation, 2024) accessed: 2026-03-23.
- [174] J. Thaler, Proofs, arguments, and zero-knowledge, [Foundations and Trends in Privacy and Security](#) **4**, 117 (2022).
- [175] N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer, From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again, in [Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCs '12](#) (Association for Computing Machinery, New York, NY, USA, 2012) pp. 326–349.
- [176] D. Koshelev, [Generation of "independent" points on elliptic curves by means of mordell–weil lattices](#), Cryptology ePrint Archive, Paper 2022/794 (2022).
- [177] Wikipedia contributors, [Nothing-up-my-sleeve number](#) (2026).
- [178] D. R. L. Brown, [Rolling up sleeves when subversion’s in a field?](#), Cryptology ePrint Archive, Paper 2020/074 (2020).
- [179] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell, [Bulletproofs: Short proofs for confidential transactions and more](#), Cryptology ePrint Archive, Paper 2017/1066 (2017).
- [180] L. Gudgeon, P. Moreno-Sanchez, S. Roos, P. McCorry, and A. Gervais, [SoK: Layer-two blockchain protocols](#), Cryptology ePrint Archive, Paper 2019/360 (2019).
- [181] Zcash Team, [The Pasta curves: Pallas and Vesta](#), GitHub Repository (2026), technical specification and implementation of a cycle of prime-order elliptic curves designed for highly efficient recursive proof composition in Halo 2.
- [182] evm.codes Contributors, [Precompiled contracts](#), Interactive EVM Reference (Osaka Hardfork Spec) (2026), catalog of precompiled contracts on Ethereum, including ECDSA recovery (0x01) and BLS12-381 operations (0x0b-0x11). Source code available at: <https://github.com/duneanalytics/evm.codes>.
- [183] P. Saggese, M. Fröwis, S. Kitzler, B. Haslhofer, and R. Auer, [Towards verifiability of total value locked \(TVL\) in decentralized finance](#), BIS Working Papers 1268 (Bank for International Settlements, 2025) available at: <https://www.bis.org/publ/work1268.pdf>.
- [184] J. Poon and T. Dryja, [The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments](#), Protocol White Paper DRAFT Version: 0.5.9.2 (Lightning Network, 2016).
- [185] V. Buterin, [Ethereum: A Next-Generation Smart Contract & Decentralized Application Platform](#), Protocol White Paper (Ethereum Foundation, 2014).
- [186] G. Wood, [Ethereum: A Secure Decentralised Generalised Transaction Ledger](#), Formal Specification Version: efc5f9a (Ethereum Foundation, 2025) shanghai Version. Accessed: 2026-03-22.
- [187] N. Szabo, Smart contracts: Building blocks for digital markets (2018).
- [188] Y. Zhao, [The Cost of Delay: Evidence from the Ethereum Transaction Fee Market](#), [Economics honors thesis](#), Duke University, Durham, NC (2023), accessed: 2026-03-21.
- [189] M. Kalinin, D. Ryan, and V. Buterin, [EIP-3675: Upgrade consensus to Proof-of-Stake](#), Ethereum Improvement Proposal 3675 (Ethereum Foundation, 2021) accessed: 2026-03-23.
- [190] V. Buterin, D. Hernandez, T. Kampefner, K. Pham, Z. Qiao, D. Ryan, J. Sin, Y. Wang, and Y. X. Zhang, Combining ghost and casper, [arXiv:2003.03052](#) (2020).
- [191] A. Lewis-Pye and T. Roughgarden, Permissionless consensus, [arXiv:2304.14701](#) (2023).
- [192] D. Johnson, A. Menezes, and S. Vanstone, The elliptic curve digital signature algorithm (ecdsa), [International Journal of Information Security](#) **1**, 36 (2001).
- [193] S. Palladino, F. Giordano, and H. Croubois, [ERC-1967: Proxy Storage Slots](#), Ethereum Improvement Proposal 1967 (Ethereum Foundation, 2019) accessed: 2026-03-23.
- [194] V. Buterin, Y. Weiss, D. Tirosh, S. Nacson, A. Forshtat, K. Gazso, and T. Hess, [ERC-4337: Account Abstraction Using Alt Mempool \[DRAFT\]](#), Ethereum Improvement Proposal 4337 (Ethereum Foundation, 2021) accessed: 2026-03-23.
- [195] A. Day and E. Medvedev, [Ethereum in BigQuery: a public dataset for smart contract analytics](#), Google Cloud Blog (2018), developed in collaboration with Kaggle, AMPOS, and the Ethereum ETL project.

- [196] V. Buterin, S. Wilson, A. Dietrichs, and lightclient, *EIP-7702: Set Code for EOAs*, Ethereum Improvement Proposal 7702 (Ethereum Foundation, 2024) accessed: 2026-03-23.
- [197] F. Vogelsteller and V. Buterin, *ERC-20: Token Standard*, Ethereum Improvement Proposal 20 (Ethereum Foundation, 2015) accessed: 2026-03-23.
- [198] W. Entriiken, D. Shirley, J. Evans, and N. Sachs, *ERC-721: Non-Fungible Token Standard*, Ethereum Improvement Proposal 721 (Ethereum Foundation, 2018) accessed: 2026-03-23.
- [199] J. Lebrun, T. Malghem, K. Thizy, L. Falempin, and A. Boudjemaa, *ERC-3643: T-REX - Token for Regulated EXchanges*, Ethereum Improvement Proposal 3643 (Ethereum Foundation, 2021) accessed: 2026-03-23.
- [200] L. Fink and R. Goldstein, Larry Fink and Rob Goldstein on how tokenisation could transform finance, *The Economist By Invitation (2025)*, reprint available at: <https://www.blackrock.com/corporate/literature/article-reprint/larry-fink-rob-goldstein-economist-op-ed-tokenization.pdf>.
- [201] Castle Labs, Inc., Real-World Asset Analytics Dashboard (RWA.xyz), <https://app.rwa.xyz> (2026), accessed: 2026-03-21.
- [202] A. Zamyatin, M. Al-Bassam, D. Zindros, E. Kokoris-Kogias, P. Moreno-Sanchez, A. Kiayias, and W. J. Knottenbelt, Sok: Communication across distributed ledgers, in *Financial Cryptography and Data Security: 25th International Conference, FC 2021, Virtual Event, March 1–5, 2021, Revised Selected Papers, Part II* (Springer-Verlag, Berlin, Heidelberg, 2021) pp. 3–36.
- [203] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev, Scalable zero knowledge with no trusted setup, in *Advances in Cryptology–CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part III* (Springer-Verlag, Berlin, Heidelberg, 2019) pp. 701–732.
- [204] J. Kempton, *EIP-7932: Secondary Signature Algorithms [DRAFT]*, Ethereum Improvement Proposal 7932 (Ethereum Foundation, 2025) accessed: 2026-03-23.
- [205] L2BEAT sp. z o.o., *L2beat: The state of the layer two ecosystem* (2026).
- [206] V. Buterin and V. Griffith, Casper the friendly finality gadget, [arXiv:1710.09437](https://arxiv.org/abs/1710.09437) (2017).
- [207] F. D’Amato, B. Monnot, M. Neuder, Potuz, J. Traglia, and T. Tsao, *EIP-7732: Enshrined Proposer-Builder Separation [DRAFT]*, Ethereum Improvement Proposal 7732 (Ethereum Foundation, 2024) accessed: 2026-03-23.
- [208] hildobby, *Ethereum ETH staking dashboard*, Dune Analytics (2026), industry-standard real-time tracker for Ethereum staking metrics.
- [209] J. Drake, D. Khovratovich, M. Kudinov, and B. Wagner, *Hash-based multi-signatures for post-quantum ethereum*, Cryptology ePrint Archive, Paper 2025/055 (2025).
- [210] V. Buterin, D. Feist, D. Loerakker, G. Kadianakis, M. Garnett, M. Taiwo, and A. Dietrichs, *EIP-4844: Shard Blob Transactions*, Ethereum Improvement Proposal 4844 (Ethereum Foundation, 2022) accessed: 2026-03-23.
- [211] Ethereum Foundation, *The KZG trusted setup ceremony*, Official Protocol Website (2022), a massively multi-party computation to produce the Structured Reference String for KZG polynomial commitments used in EIP-4844.
- [212] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev, Fast reed-solomon interactive oracle proofs of proximity (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018) pp. 14:1–14:17.
- [213] M. Hall-Andersen, M. Simkin, and B. Wagner, *Foundations of data availability sampling*, Cryptology ePrint Archive, Paper 2023/1079 (2023).
- [214] Ethereum Foundation, Official Website of the Ethereum Foundation, <https://ethereum.foundation/> (2026), accessed: 2026-03-22.
- [215] A. Antonopoulos and G. D., *Mastering Ethereum: Building Smart Contracts and DApps*, 1st ed. (O’Reilly Media, 2018).
- [216] C. Detrio, *EIP-779: Hardfork Meta: DAO Fork*, Ethereum Improvement Proposal 779 (Ethereum Foundation, 2017) accessed: 2026-03-23.
- [217] U.S. Securities and Exchange Commission, *Report of Investigation Pursuant to Section 21(a) of the Securities Exchange Act of 1934: The DAO*, Report of Investigation Release No. 81207 (U.S. Securities and Exchange Commission, 2017).
- [218] A. Schoedon, *EIP-999: Restore Contract Code at 0x863DF6BFa4469f3ead0bE8f9F2AAE51c91A907b4 [DRAFT]*, Ethereum Improvement Proposal 999 (Ethereum Foundation, 2018) accessed: 2026-03-23.
- [219] j. Drake, *Today marks an inflection in the ethereum foundation’s long-term quantum strategy.*, X (formerly Twitter) (20256).
- [220] L. Parker, *Stepping back*, GitHub Gist (2024), technical manifesto calling for a moratorium on quantum-vulnerable R&D and the immediate development of a fallback post-quantum Monero protocol.
- [221] SyntheticBird, *Discussion: Post-quantum security and ethical considerations over elliptic curve cryptography*, Monero Research Lab (MRL), GitHub Issue #131 (2024), formalizes the proposal for a 5-year post-quantum migration timeline, citing ethical responsibilities and the Google Willow quantum processor announcement.
- [222] J. Jefferson, *DIP: OP.CHECKZKP zero-knowledge upgrade*, Dogecoin Improvement Proposal (Draft), GitHub Issue #3869 (2025), proposes a new opcode for onchain verification of SNARKs to enable zk-rollups on Dogecoin.
- [223] A. Back, M. Corallo, L. Dashjr, M. Friedenbach, G. Maxwell, A. Miller, A. Poelstra, J. Tim’on, and P. Wuille, *Enabling Blockchain Innovations with Pegged Sidechains*, Protocol White Paper Commit: 5620e43 (Blockstream, 2014).
- [224] A. Singh, K. Click, R. M. Parizi, Q. Zhang, A. Dehghantanha, and K.-K. R. Choo, Sidechain technologies in blockchain networks: An examination and state-of-the-art review, *J. Netw. Comput. Appl.* **149**, 10.1016/j.jnca.2019.102471 (2020).
- [225] S. D. Lerner, J. A. Cid-Fuentes, J. Len, R. Fernández-València, P. Gallardo, N. Vescovo, R. Laprida, S. Mishra, F. Jinich, and D. Masini, *RSK: A bitcoin sidechain with stateful smart-contracts*, Cryptology ePrint Archive, Paper 2022/684 (2022).
- [226] S. D. Lerner, *RSK: Bitcoin Powered Smart Contracts*, Protocol White Paper Revision 11 (RSK Labs, 2019).
- [227] Tom Elvis Jedusor, *MIMBLEWIMBLE*, Original text file distributed via IRC (2016).

- [228] A. Yang, D. Burkett, and C. Lee, *LIP 003: MumbleWimble via Extension Blocks (Consensus layer)*, Litecoin Improvement Proposal 3 (Litecoin Project, 2019) accessed: 2026-03-22.
- [229] T. Ruffing, *Switch commitments (again)*, MumbleWimble Mailing List Archive (2018), proposes to commit with $C = vG + bH$ where $b = b' + \text{hash}(vG + b'H, b'J)$, so C can be treated as either Pedersen or ElGamal commitment.
- [230] E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza, *Succinct non-interactive zero knowledge for a von Neumann architecture*, Cryptology ePrint Archive, Paper 2013/879 (2013).
- [231] D.-E. Hopwood, *ZIP 211: Disabling Addition of New Value to the Sprout Chain Value Pool*, Zcash Improvement Proposal 211 (Electric Coin Co., 2019) accessed: 2026-03-21.
- [232] Y. T. Lai and Zcash Developers, *Pull request #4489: Disable addition of new value to the Sprout value pool*, Zcash GitHub Repository (2020), implements ZIP-211.
- [233] J. Groth, *On the size of pairing-based non-interactive arguments*, Cryptology ePrint Archive, Paper 2016/260 (2016).
- [234] A. Miller and S. Bowe, *Announcing the world's largest multi-party computation ceremony*, Zcash Foundation Blog (2017), official announcement of the Powers of Tau MPC ceremony for the Sapling upgrade.
- [235] Electric Coin Co., *The halo2 Book* (2026).
- [236] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels* (1997).
- [237] D.-E. Hopwood and J. Grigg, *ZIP 2005: Quantum recoverability*, Zcash Improvement Proposal (Draft) (2025), credits: Sean Bowe. Discussion available at GitHub Issue #1135: <https://github.com/zcash/zips/issues/1135>.
- [238] S. Bowe, Zcash and Quantum Computers, <https://seanbowe.com/blog/zcash-and-quantum-computers/> (2025), accessed: 2026-03-22.
- [239] Electric Coin Company, *Turnstile enforcement against counterfeiting*, Official Zcash Blog (2019).
- [240] S. Bowe and D.-E. Hopwood, *ZIP 209: Prohibit negative shielded chain value pool balances*, Zcash Improvement Proposal (ZIP) 209 (2019).
- [241] D. A. Cooper, D. C. Apon, Q. H. Dang, M. S. Davidson, M. J. Dworkin, and C. A. Miller, *Recommendation for Stateful Hash-Based Signature Schemes* (2020).
- [242] L. Ducas, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehle, *CRYSTALS – dilithium: Digital signatures from module lattices*, Cryptology ePrint Archive, Paper 2017/633 (2017).
- [243] *Module-lattice-based digital signature standard* (2024).
- [244] R. C. Merkle, A certified digital signature, in *Advances in Cryptology — CRYPTO' 89 Proceedings*, edited by G. Brassard (Springer New York, New York, NY, 1990) pp. 218–238.
- [245] C. Dods, N. P. Smart, and M. Stam, Hash based digital signature schemes, in *Proceedings of the 10th International Conference on Cryptography and Coding*, IMA'05 (Springer-Verlag, Berlin, Heidelberg, 2005) pp. 96–115.
- [246] Algorand Foundation, *Algorand TEAL: Transaction execution approval language*, Algorand Developer Documentation (2026).
- [247] Algorand Foundation, *Algorand Virtual Machine*, Algorand Developer Documentation (2026).
- [248] R. Perlner, FIPS 206 Status Update, Technical Presentation, [https://csrc.nist.gov/csrc/media/presentations/2025/fips-206-fn-dsa-\(falcon\)/images-media/fips_206-perlner_2.1.pdf](https://csrc.nist.gov/csrc/media/presentations/2025/fips-206-fn-dsa-(falcon)/images-media/fips_206-perlner_2.1.pdf) (2025), presented at the 6th PQC Standardization Conference in Gaithersburg, Maryland.
- [249] P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang, *Falcon: Fast-Fourier Lattice-based Compact Signatures over NTRU*, Technical Specification (NIST PQC Submission) Specification v1.2 (NIST Post-Quantum Cryptography Project, 2020).
- [250] Algorand Foundation, *Algorand TEAL opcodes: falcon.verify*, Algorand Developer Reference (2026), bytecode 0x85, available in v12.
- [251] Algorand Foundation, *Rekeying accounts*, Algorand Developer Documentation (2026).
- [252] D. Little, Silur, and I. E. Ashimine, *Solana Winternitz Vault: A quantum-resistant lamport vault using WOTS*, GitHub Repository (2026).
- [253] M. Swayne, *Solana takes a step toward PQC era with quantum-resistant vault*, The Quantum Insider (2025).
- [254] Castle Labs, Inc., OpenEden T-Bill (TBILL) Vault — Real-World Asset Analytics Dashboard (RWA.xyz), <https://app.rwa.xyz/assets/TBILL> (2026), accessed on 2026-02-24 and observed 61.8 million USD on the XRP Ledger out of total supply of 92.5 million USD.
- [255] C. Goes, The interblockchain communication protocol: An overview, [arXiv:2006.15918](https://arxiv.org/abs/2006.15918) (2006).
- [256] Circle Internet Financial, *Circle is discontinuing support for USDC on the TRON blockchain*, Official Blog Announcement (2024).
- [257] The QRL Foundation, *QRL 2.0 (Codename Project Zond)*, Official Project Roadmap (2026).
- [258] M. M. T. Chakravarty, J. Chapman, K. MacKenzie, O. Melkonian, M. Peyton Jones, and P. Wadler, The extended utxo model, in *Financial Cryptography and Data Security*, edited by M. Bernhard, A. Bracciali, L. J. Camp, S. Matsuo, A. Maurushat, P. B. Rønne, and M. Sala (Springer International Publishing, Cham, 2020) pp. 525–539.
- [259] V. L. s.r.o., *Understanding Cardano certificates*, Cexplorer.io (2023), technical analysis of Cardano's stake pool and delegation certificates, detailing the separation of payment and staking credentials.
- [260] TRON DAO, *Account permission management*, TRON Developer Network (2026).
- [261] Ripple, *Documentation – best practices – key management – remove a regular key pair*, XRP Ledger Dev Portal (2026).
- [262] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM Journal on Computing* **26**, 1484 (1997).
- [263] G. Alagic, L. Chen, S. Jordan, Y.-K. Liu, D. Moody, R. Peralta, R. Perlner, and D. Sonneman, *Transition to Post-Quantum Cryptography Standards*, NIST Internal Report (Initial Public Draft) NIST IR 8547 (National Institute of

Standards and Technology, 2024).

- [264] L. Lamport, *Constructing Digital Signatures from a One Way Function*, Tech. Rep. CSL-98 (1979) this paper was published by IEEE in the Proceedings of HICSS-43 in January, 2010.
- [265] L. Chen, S. Jordan, Y.-K. Liu, D. Moody, R. Peralta, R. Perlner, and D. Smith-Tone, *Report on Post-Quantum Cryptography* (2016).
- [266] C. Peikert, *A decade of lattice cryptography*, Cryptology ePrint Archive, Paper 2015/939 (2015).
- [267] E. Fathalla and M. Azab, Beyond classical cryptography: A systematic review of post-quantum hash-based signature schemes, security, and optimizations, *IEEE Access* **12**, 175969 (2024).
- [268] V. Weger, N. Gassner, and J. Rosenthal, A survey on code-based cryptography, [arXiv:2201.07119](https://arxiv.org/abs/2201.07119) (2022).
- [269] J. Ding, J. Gower, and D. Schmidt, *Multivariate Public Key Cryptosystems*, Vol. 25 (2006).
- [270] L. De Feo, Mathematics of isogeny based cryptography, [arXiv:1711.04062](https://arxiv.org/abs/1711.04062) (2017).
- [271] D. Jao and L. De Feo, Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies, in *Post-Quantum Cryptography* (Springer Berlin Heidelberg, 2011) pp. 19–34.
- [272] W. Castryck and T. Decru, *An efficient key recovery attack on SIDH*, Cryptology ePrint Archive, Paper 2022/975 (2022).
- [273] O. Regev, On lattices, learning with errors, random linear codes, and cryptography, in *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '05 (Association for Computing Machinery, New York, NY, USA, 2005) pp. 84–93.
- [274] S. P. Jordan, N. Shuttly, M. Wootters, A. Zalcman, A. Schmidhuber, R. King, S. V. Isakov, T. Khattar, and R. Babbush, Optimization by decoded quantum interferometry, *Nature* **646**, 831 (2025).
- [275] M. Ajtai, Generating hard instances of the short basis problem, in *Automata, Languages and Programming*, edited by J. Wiedermann, P. van Emde Boas, and M. Nielsen (Springer Berlin Heidelberg, Berlin, Heidelberg, 1999) pp. 1–9.
- [276] A. Wein, A. El Alaoui, and C. Moore, The kikuchi hierarchy and tensor pca, *J. ACM* **72**, 10.1145/3762806 (2025).
- [277] M. B. Hastings, Classical and quantum algorithms for tensor principal component analysis, *Quantum* **4**, 237 (2020).
- [278] A. Schmidhuber, R. O'Donnell, R. Kothari, and R. Babbush, Quartic quantum speedups for planted inference, *Physical Review X* **15**, 10.1103/physrevx.15.021077 (2025).
- [279] *Module-lattice-based key-encapsulation mechanism standard* (2024).
- [280] *Stateless hash-based digital signature standard* (2024).
- [281] S. Marzougui, I. Kabin, J. Krämer, T. Aulbach, and J.-P. Seifert, *On the feasibility of single-trace attacks on the gaussian sampler using a CDT*, Cryptology ePrint Archive, Paper 2023/142 (2023).
- [282] PQShield, *Post-quantum signatures zoo*, GitHub Pages (2026).
- [283] L. D. Feo, D. Kohel, A. Leroux, C. Petit, and B. Wesolowski, *SQISign: compact post-quantum signatures from quaternions and isogenies*, Cryptology ePrint Archive, Paper 2020/1240 (2020).
- [284] H. Beast, *P2QRH / BIP-360 update*, Bitcoin Development Mailing List (Google Groups) (2025), discussion of BIP-360 featuring contributions from Matt Corallo, Jonas Nick, and others. Highlights the DDoS risk of Post-Quantum transition scripts in Bitcoin.
- [285] P. C. van Oorschot and M. J. Wiener, Parallel collision search with cryptanalytic applications, *Journal of Cryptology* **12**, 1 (1999).
- [286] N. Bindel and B. Hale, *A note on hybrid signature schemes*, Cryptology ePrint Archive, Paper 2023/423 (2023).
- [287] M. A. Aardal, D. F. Aranha, K. Boudgoust, S. Kolby, and A. Takahashi, *Aggregating falcon signatures with LaBRADOR*, Cryptology ePrint Archive, Paper 2024/311 (2024).
- [288] Y. Quan, Improving bitcoin's post-quantum transaction efficiency with a novel lattice-based aggregate signature scheme based on crystals-dilithium and a stark protocol, *IEEE Access* **10**, 132472 (2022).
- [289] N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer, *Recursive composition and bootstrapping for SNARKs and proof-carrying data*, Cryptology ePrint Archive, Paper 2012/095 (2012).
- [290] I. Stewart, D. Ilie, A. Zamyatin, S. Werner, M. F. Torshizi, and W. J. Knottenbelt, Committing to quantum resistance: a slow defence for bitcoin against a fast quantum computing attack, *Royal Society Open Science* **5**, 180410 (2018).
- [291] N. Sklavos and O. Koufopavlou, Implementation of the sha-2 hash family standard using fpgas, *The Journal of Supercomputing* **31**, 227 (2005).
- [292] S. D. Lerner, *The well deserved fortune of Satoshi Nakamoto, Bitcoin creator, visionary and genius*, Bitslog (2013), seminal research identifying the "Patoshi" mining pattern and estimating the size of Satoshi Nakamoto's early Bitcoin holdings.
- [293] M. Humayun and R. Belk, "satoshi is dead. long live satoshi": The curious case of bitcoin's creator (2018) pp. 19–35.
- [294] K. Finlow-Bates, *In search of lost keys*, Medium (2023).
- [295] P. McCormack and D. Held, *Bitcoin's immaculate conception: Audio interview transcription — WBD043*, Medium: HackerNoon (2018), interview Date: Oct 25, 2018. Discusses the philosophical significance of Satoshi Nakamoto's disappearance and the "Immaculate Conception" theory of Bitcoin's launch.
- [296] Officer's Notes, *New theory by BTCparser: Who might be the real Satoshi? answering the question of why he never touched his 2009 Bitcoins!*, Medium: Coinmonks (2024).
- [297] D. Bradbury, The problem with bitcoin, *Computer Fraud & Security* **2013**, 5 (2013).
- [298] J. Bartusek, *Obfuscation of Quantum Computation*, Ph.D. thesis, EECS Department, University of California, Berkeley (2024).
- [299] M.-Y. Huang and E.-C. Tang, Obfuscation of unitary quantum programs, [arXiv:2507.11970](https://arxiv.org/abs/2507.11970) (2025).
- [300] Martini, Luca and Heuser, Matthias and Fest, Martin and Stegemann, Uwe and Schneider, Sebastian and Brenna, Gabriel and Windhagen, Eckart and Poppensieker, Thomas, *Bad Banks: Finding the Right Exit from the Financial Crisis*, Working

Paper 12 (McKinsey Working Papers on Risk, New York, NY, 2009).

- [301] H. Beast and M. Casey, **BIP Draft: Hourglass spending rules**, Bitcoin Improvement Proposal (Draft), GitHub Repository (2025), proposes a soft fork consensus rule restricting P2PK throughput to mitigate inflationary quantum attacks.
- [302] M. Palatinus, P. Rusnak, A. Voisine, and S. Bowe, **BIP 39: Mnemonic code for generating deterministic keys**, Bitcoin Improvement Proposal 39 (Bitcoin Core, 2013) source: <https://github.com/bitcoin/bips/blob/master/bip-0039.mediawiki>. Accessed: 2026-03-22.
- [303] C. Deegan, **hello yellowpages**, Project Eleven Blog (2025).
- [304] J. Nick, A. Poelstra, and G. Sanders, **Liquid: A Bitcoin Sidechain**, Protocol White Paper (Blockstream, 2020).
- [305] E. Voskuil, **Cryptoeconomics: Fundamental Principles of Bitcoin**, 2nd ed. (Eric Voskuil, United States, 2020) version 1.2.3 (PDF). Edited and illustrated by James Chiang.
- [306] M. Nadler and F. Schär, Tornado cash and blockchain privacy: A primer for economists and policymakers, **Federal Reserve Bank of St. Louis Review** (2023), first Quarter 2023.
- [307] A. Wahrst"atter, J. Ernstberger, A. Yaish, L. Zhou, K. Qin, T. Tsuchiya, S. Steinhorst, D. Svetinovic, N. Christin, M. Barzentewicz, and A. Gervais, Blockchain censorship, in **Proceedings of the ACM Web Conference 2024**, WWW '24 (Association for Computing Machinery, New York, NY, USA, 2024) pp. 1632–1643.
- [308] A. Brownworth, J. Durfee, M. J. Lee, and A. Martin, **Regulating Decentralized Systems: Evidence from Sanctions on Tornado Cash**, Staff Reports 1112 (Federal Reserve Bank of New York, 2024).
- [309] N. Tovanich, N. Soulie, N. Heulot, and P. Isenberg, The evolution of mining pools and miners' behaviors in the bitcoin blockchain, **IEEE Transactions on Network and Service Management** **19**, 3633 (2022).
- [310] S. Shen and A. Galbraith, **China's ban forces some bitcoin miners to flee overseas, others sell out**, Reuters (2021).
- [311] Reuters, **US becomes largest Bitcoin mining centre following China ban**, Reuters (2021).
- [312] Cambridge Centre for Alternative Finance, **Cambridge Bitcoin Electricity Consumption Index (CBECI): Mining map**, Interactive Data Dashboard (2026), approximate geographic distribution of Bitcoin hashrate and estimated total network electricity consumption.
- [313] G. Walker, Taproot: Technical Explanation — Learn Me A Bitcoin, <https://learnmeabitcoin.com/technical/upgrades/taproot/> (2026), accessed: 2026-03-21.
- [314] C. S. Gokhale and A. Traulsen, Evolutionary games in the multiverse, **Proceedings of the National Academy of Sciences** **107**, 5500 (2010), <https://www.pnas.org/doi/pdf/10.1073/pnas.0912214107>.
- [315] W. Poundstone, **Prisoner's Dilemma** (Doubleday, 1992).
- [316] G. H. Snyder, "prisoner's dilemma" and "chicken" models in international politics, **International Studies Quarterly** **15**, 66 (1971), <https://academic.oup.com/isq/article-pdf/15/1/66/5096260/15-1-66.pdf>.
- [317] R. O. Keohane, **After Hegemony: Cooperation and Discord in the World Political Economy**, rev - revised ed. (Princeton University Press, 1984).
- [318] W. Kenton, **Escheat: Meaning, process, and reclaiming assets**, Investopedia (2025), reviewed by Toby Walters, defines "escheat" as "when a government obtains ownership of unclaimed property or estate assets due to there being no identifiable heirs or beneficiaries".
- [319] Merriam-Webster, **Definition of Laches**, Merriam-Webster.com Dictionary, defines "laches" as "undue delay in asserting a legal right or privilege".
- [320] Uniform Law Commission, **Revised Uniform Unclaimed Property Act**, Uniform Law Commission (2016), last updated in 2016; originally promulgated in 1954.
- [321] G. Walker, **OP_RETURN: Storing data on the blockchain — Learn Me A Bitcoin**, <https://learnmeabitcoin.com/technical/script/return/> (2025), accessed: 2026-03-21.
- [322] L. Lessig, **Code and Other Laws of Cyberspace** (Basic Books, New York, NY, 1999).
- [323] A. L. Russell, 'Rough Consensus and Running Code' and the Internet-OSI Standards War, **IEEE Ann. Hist. Comput.** **28**, 48 (2006).
- [324] C. Gidney, Verifying Measurement Based Uncomputation — Algassert, <https://algassert.com/post/1903> (2019), accessed: 2026-03-22.
- [325] P. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in **Proceedings 35th Annual Symposium on Foundations of Computer Science** (1994) pp. 124–134.
- [326] M. Ekerå, Revisiting shor's quantum algorithm for computing general discrete logarithms, **arXiv:1905.09084** (2019).
- [327] R. B. Griffiths and C.-S. Niu, Semiclassical Fourier Transform for quantum computation, **Physical Review Letters** **76**, 3228 (1996).
- [328] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, Qubitization of arbitrary basis quantum chemistry leveraging sparsity and low rank factorization, **Quantum** **3**, 208 (2019).
- [329] C. Gidney, Halving the cost of quantum addition, **Quantum** **2**, 74 (2018).
- [330] C. Zalka, **Shor's algorithm with fewer (pure) qubits** (2006).
- [331] **SHA-3 standard: permutation-based hash and extendable-output functions** (2015).
- [332] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev, **Scalable, transparent, and post-quantum secure computational integrity**, Cryptology ePrint Archive, Paper 2018/046 (2018).
- [333] P. Valiant, Incrementally verifiable computation or proofs of knowledge imply time/space efficiency, in **Theory of Cryptography**, edited by R. Canetti (Springer Berlin Heidelberg, Berlin, Heidelberg, 2008) pp. 1–18.
- [334] L. Goldberg, S. Papini, and M. Riabzev, **Cairo—a turing-complete STARK-friendly CPU architecture**, Cryptology ePrint Archive, Paper 2021/1063 (2021).

Given the costs of our point addition circuit, the optimal window size is $w = 16$. At this window size, a total of

$$\left(\frac{2 \times 256}{16} - 4\right) = 28 \tag{A3}$$

windowed point additions are performed for the 256-bit ECDLP algorithm.

Compared to the millions of gates used by the windowed lookup point additions, there are negligible additional costs in the final circuit. For example, we use measurement based uncomputation to uncompute the table lookups [328] which has a small cost associated with it. Similarly, each phase estimation must be accompanied by a small amount of work associated with performing a qubit-recycled Quantum Fourier Transform [327]. We note again that these details are not expected to materially help an adversary as they have been common ingredients of all papers on using a quantum computer to break ECDLP since 2019.

4. Measurement Based Uncomputation and Kickmix Circuits

We use measurement based uncomputation (MBUC) [93, 324, 329] to clear intermediate ancilla qubits. Traditional reversible computation requires executing the inverse circuit U^\dagger to cleanly disentangle the intermediate ancilla qubits used as temporary scratchpad, effectively doubling the logical gate counts. MBUC offers a cheaper, measurement driven alternative. Again, this is a common thread in all recent work and so the fact we use MBUC should not be surprising to an adversary. Suppose, in every branch of the superposition, an ancilla qubit q holds a deterministic boolean function $f(k)$ of the data register, represented by the state $\sum \alpha_k |k\rangle |f(k)\rangle$. Measuring q in the Pauli X basis yields two equally probable outcomes:

- $|+\rangle$ (False): The qubit q is successfully disentangled and erased. The system collapses to $\sum \alpha_k |k\rangle$. No additional uncomputation operations are needed.
- $|-\rangle$ (True): The qubit q is erased but a phase kickback occurs, negating the amplitudes of the states where $f(k) = 1$. The system collapses to $\sum \alpha_k (-1)^{f(k)} |k\rangle$.

In the latter case, the uncomputation is incomplete until a phase correction is applied. This is possible because the data register $|k\rangle$ is preserved and the phase correction is a function of k . The cost of this phase correction is usually slightly smaller [329], and sometimes much smaller [328], than the cost of computing $f(k)$.

We define kickmix circuits as the specific subclass of quantum circuits composed entirely of classical reversible logic gates, measurement-based uncomputation via Pauli-X basis measurements, and diagonal phasing gates for phase correction. By restricting the architecture exclusively to this gate set, we ensure the circuit never generates intractable quantum entanglement and can be efficiently simulated classically [324]. The etymology of *kickmix* derives from the circuit’s ability to support phase *kickback* and *mixing* classical states.

5. Verifiable Fuzz Testing via the Fiat-Shamir Heuristic

Shor’s algorithm functions well even if a small percentage of values in the superposition are wrong — a superposition with 1% of points in the wrong place will cause the algorithm to fail at most 1% of the time [46, 86, 330]. Consequently, it is sufficient for us to prove our circuits are approximately correct rather than exactly correct.

For the proof of approximate correctness, we use fuzz testing. We prove 9024 random inputs, chosen by the Fiat-Shamir heuristic [91], are mapped to the correct output. This is sufficient for 128 bits of cryptographic security that the circuit is acting correctly on at least 99% of inputs, because if the circuit mapped more than 1% of inputs to the wrong output then the probability of it passing 9024 independent random tests without a single failure is at most $(1 - 0.01)^{9024} \approx 2^{-130}$. To implement the Fiat-Shamir heuristic, we use the SHAKE256 extendable-output function (XOF) [331] seeded with the bytes of the secret circuit as a Cryptographically Secure Pseudo Random Number Generator (CSPRNG). Each test point is generated as $G \cdot k$ where G is the curve’s generator and k is a random 256 bit integer sampled using this CSPRNG.

6. Cryptographic Attestation via SP1 and Groth16 SNARK

To provide a publicly verifiable proof that this Fiat-Shamir fuzz testing was executed honestly on a circuit C satisfying our claimed resource bounds, without publishing C itself, we simulate C inside the SP1 Zero-Knowledge

Virtual Machine (zkVM) [83]. SP1 acts as a standard RISC-V processor that mathematically arithmetizes its own execution. We implement a guest program to be run inside the SP1 zkVM consisting of our kickmix circuit simulator, the Fiat-Shamir testing framework, and the secp256k1 elliptic curve point addition logic in standard Rust, which SP1 compiles directly into RISC-V bytecode.

Inside the zkVM, the guest program performs four critical tasks. First, it computes a SHA-256 hash of the input circuit and commits it as a public output, along with the demanded resource counts that the input circuit must satisfy. Second, it uses the SHAKE256 extendable-output function (XOF) [331] initialized with the bytes of the secret circuit to generate 9024 test inputs via the Fiat-Shamir transform, as described in the section above. Third, it simulates the parsed circuit using the kickmix simulator on the 9024 Fiat-Shamir derived test cases to prove functional correctness of the parsed circuit. The kickmix simulator uses the XOF initialized above for input test generation, as a pseudo-random number generator (PRNG) to simulate quantum measurements in MBUC. The simulator also records the total number of non-Clifford gates (CCX+CCZ) executed by the circuit across the 9024 test runs. Finally, the program asserts that the input private circuit satisfies the demanded resource counts — verifying that $C_{\text{low-qubit}}$ executes at most 2,700,000 Toffoli gates and requires at most 1175 logical qubits and $C_{\text{low-gate}}$ executes at most 2,100,000 Toffoli gates and requires at most 1425 logical qubits.

As SP1 simulates the guest program, it records the state of the CPU at every clock cycle into an execution trace matrix. To cryptographically attest to this multi-billion-cycle execution, SP1 divides the computation into shards of roughly 2^{22} RISC V instructions, generates a STARK proof [332] of each shard, and leverages recursive proof composition [333, 334] to generate a final compressed STARK proof. However, since individual STARK proofs in SP1 do not currently satisfy the Zero Knowledge Property [335], we invoke the SP1 prover in Groth16 mode which wraps the compressed hash-based STARK proof into a Groth16 zk-SNARK [233] and ensures Zero Knowledge and Succinctness. The resulting zero-knowledge proof artifact is published alongside our prover and verifier code in the accompanying Zenodo dataset [92].

Finally, we note a practical irony in our choice of cryptographic attestation. Because the Groth16 SNARK relies on pairing-friendly elliptic curves, its soundness is inherently vulnerable to the very quantum attacks analyzed in this work. In principle, one could use a sufficiently large, fault-tolerant quantum computer to forge this zero-knowledge proof. However, since cryptographically relevant quantum computers do not exist as of today, the soundness of our proof remains intact.

7. Code Availability for Zero-Knowledge Proof

The source code for generating the Zero-Knowledge (ZK) proofs, along with all associated artifacts, is archived and publicly available in our Zenodo repository [92].