The Superlinearity Problem in Post-Quantum Blockchains

Sunoo Park
Columbia University

Nicholas Spooner
University of Warwick

Abstract

The *proof of work* mechanism by which many blockchain-based protocols achieve consensus may be undermined by the use of quantum computing in mining — even when all cryptographic primitives are replaced with post-quantum secure alternatives. First, we offer an impossibility result: we prove that quantum (Grover) speedups in solving a large, natural class of proof-of-work puzzles cause an inevitable incentive incompatibility in mining, by distorting the reward structure of mining in proof-of-work-based protocols such as Bitcoin. We refer to such distortion as the Superlinearity Problem. Our impossibility result suggests that for robust post-quantum proof-of-work-based consensus, we may need to look beyond standard cryptographic models. We thus propose a proof-of-work design in a random-beacon model, which is tailored to bypass the earlier impossibility. We conclude with a discussion of open problems, and of the challenges of integrating our new proof-of-work scheme into decentralised consensus protocols under realistic conditions.

1 Introduction

Blockchain-based technologies have gained remarkable traction since the proposal of the Bitcoin protocol in 2009 [Nak09]. Today, blockchains and cryptocurrencies are familiar topics in mainstream media and among government regulators (e.g., [GYB22, Kle22, Cou, BMW22]), and the top two cryptocurrencies' collective market cap is around 500 billion U.S. dollars [Coi22].

At the same time, progress in quantum computing has been rapidly advancing. Recent experiments have shown that quantum computers can perform certain (contrived) tasks faster than the largest classical supercomputers available [AAB+19]. Much has been written on the potential impact of quantum computing on blockchain-based technologies [ABL+18, Bol20, FF20], primarily focused on the fact that existing blockchains tend to rely on *pre-quantum* cryptography (primarily signature schemes) that is breakable using sufficiently powerful quantum computers [LaM21].

Certainly, today's blockchains would become insecure if a sufficiently powerful quantum computer were developed. This will not happen overnight: quantum computers that can break today's pre-quantum cryptography are widely believed to be more than a decade away [MP21]. Moreover, this is an aspect of a much broader phenomenon, relevant not only to blockchains but to all of the essential Internet infrastructure that relies on pre-quantum cryptography, including protocols such as HTTPS and SSL that are used for viewing most websites, and to secure online banking and shopping. As such, secure and efficient *post-quantum* alternatives to such pre-quantum cryptography are already well studied, and implementation and standardization processes are well underway [Nat22]. These processes are designed to ensure that by the time that quantum computers become viable, any infrastructure that relies on pre-quantum cryptography will have replaced it with secure, standardized post-quantum alternatives.

However, quantum computing poses another lesser studied but potentially more impactful threat to many existing blockchain technologies. The *proof of work* mechanism by which many blockchain-based protocols achieve the crucial property of consensus may be undermined by protocol participants' (miners') use of quantum computers. In a nutshell, a proof of work is a *moderately hard* computational puzzle, which many protocol participants (*miners*) attempt until someone "wins" by finding a solution to the latest puzzle for a given blockchain. The "winning" miner then appends some information (a *block*) to the blockchain.

Proofs of work rely on cryptographic hash functions, which are believed to already be post-quantum secure, so they would remain moderately hard (as intended) for quantum computers. However, the *relative* power of different miners in a blockchain network would change impactfully if some or all of them had quantum computers. This causes two key problems.

I. Quantum Advantage Problem (Nearer Future). Efficient quantum computers would speed up quantum miners' production of proofs of work compared to classical miners,[1] likely discouraging those without quantum computers from participating in mining. However, this would not be problematic if and when quantum computers become widely available.

II. Quantum Superlinearity Problem (Farther Future). Because the quantum speedup is *quadratic*, more computationally powerful quantum miners would gain a disproportionate speedup, eliminating the incentive for less powerful quantum miners — as well as those who lack quantum computers entirely — to participate at all. The result could be a destablising concentration of network control among just the most computationally powerful miners in proof-of-work-based blockchains, weakening the networks' security and consensus properties, as well as undermining the vision of fairness and distributed governance that motivate many blockchain-based systems today. As a concrete example, the famous "51% attack" on Bitcoin — thus named because it requires control of 51% of network hash power — would become possible through control of just over a quarter of hash power, under certain conditions.[2]

The Quantum Superlinearity Problem has a natural classical variant, which we call the Classical Superlinearity Problem: namely, similar problems arise when more powerful *classical* miners have a disproportionate advantage over less powerful *classical* miners. Bitcoin and other major blockchains that use proofs of work are designed to yield mining advantage *proportionate* to miners' hash power. Even so, significant (and much critiqued) concentration of power has already occurred in existing proof-of-work-based blockchains due to economies of scale, specialized mining hardware, geographic disparities, and other factors [BCEM15, BMC+15, GBE+18, Shi21, LBS22], due to which some larger miners' advantage is *disproportionate* in terms of economic investment even if it is proportionate in terms of their hash power. Our results suggest that the impact of quantum superlinearity could be an order of magnitude worse than the classical counterparts, as discussed in more detail below.

We write simply Superlinearity Problem when referring to both the Quantum and Classical Superlinearity Problems.

These problems raise a natural question, which is the focus of this paper:

> *Can we design a proof of work that avoids the Quantum Advantage and Quantum Superlinearity Problems, and thus preserve the incentive structure that currently supports proof-of-work-based systems such as Bitcoin?*

We answer this question in the *negative* for a large and natural class of proofs of work encompassing all prior constructions to our knowledge, as summarized in the informal theorem

[1] *Classical* means computing without quantum computers.

[2] Classical network takeover attacks are also possible with the collusion of much less than half of mining power [ES18]. The Quantum Superlinearity Problem worsens those attacks too: basically, an attack that requires a certain fraction of classical mining power may require a much smaller fraction of quantum mining power.

below. Then, we highlight several potential directions for *positive* results outside the scope of our impossibility, and provide a partial *positive* result: a new proof-of-work construction that provably avoids the Quantum Advantage and Quantum Superlinearity problems, in a random beacon model.

Theorem (Informal). For a large, natural class of proofs of work (which includes the Hashcash [Bac02] and Equihash [BK17] methods that underlie Bitcoin and most proof-of-work blockchains today):

- the Quantum Superlinearity Problem is inherent (i.e., unavoidable), and
- the Quantum Advantage Problem is not solvable without exacerbating the Classical Superlinearity Problem.

Interpreting this impossibility in light of prior theoretical and empirical analyses of centralization in Bitcoin, it appears that the impact from widespread quantum computing could be an order of magnitude worse than the effects of superlinearity already present in the classical setting. Gencer et al. estimated recently that more than 50% of Bitcoin mining power is controlled by eight miners, and 90% is controlled by sixteen miners [GBE+18] — already a concerning centralization trend. However, Arnosti and Weinberg's theoretical model of the impact of superlinear rewards [AW22], together with our results, indicates that the equilibrium number of miners in the post-quantum setting may be *just two* for the large class of proof-of-work protocols this paper considers.

Our impossibility results suggest that to design a proof of work that avoids the Quantum Advantage and Quantum Superlinearity Problems, we may need to look beyond standard cryptographic models. We analyze our impossibility theorem in detail to highlight seven potential research directions towards post-quantum blockchains that do not suffer from the Quantum Advantage and Quantum Superlinearity Problems. Then, focusing on one of these seven directions, we propose a proof-of-work design in a random-beacon model, tailored to bypass the above impossibilities. Finally, we prove the security of our proof-of-work construction, and discuss the significant challenges that seem to remain to integrate a proof of work like ours into a realistic blockchain protocol.

Proof-of-work alternatives. Given that the Quantum Advantage and Superlinearity Problems are inherent to a large class of proofs of work, it is also natural to consider whether alternatives to proofs of work could resolve these problems. That is: *can we develop alternative consensus mechanisms not involving proofs of work, that preserve the incentive structure that currently supports proof-of-work-based systems such as Bitcoin, even in the presence of quantum computers?*

This paper's main focus is to examine what is possible and impossible within the proof-of-work approach. As such, we make just a few remarks on proof-of-work alternatives, and leave this question open as an important direction for future work. Despite significant environmental and efficiency concerns about proof-of-work-based consensus (e.g., [Vra17, Bli18, KO19, Whi22, Osb22]), it remains the dominant consensus model in blockchains today. The main competing approach of *proof of stake* [KN12, Nxt14] has not yet gained traction competitive with Bitcoin's original proof-of-work model. However, this may be changing, with the very recent shift (in early 2022) of the second-biggest cryptocurrency, Ethereum, to a proof-of-stake model [Eth, Cas22]. Yet other alternatives to proof of work exist as well, with much less adoption than proof of stake (e.g., [SSP13, Pro17, PKF+18, Pie19a, CP19, Wag]).

Existing efficient implementations of proofs of stake rely on pre-quantum cryptography much more advanced[3] than the cryptography typically used in proof-of-work blockchains. While post-quantum alternatives to these advanced cryptographic tools exist in theory, current standalone constructions would entail impractical computational overhead (e.g., in the order of 1000–10000× for certain operations [BDE+22]). An key research direction to make proof-of-stake blockchains practical for post-quantum use is to improve this overhead.

Prior work on the Quantum Advantage Problem. Past research has considered the extent to which the Quantum Advantage Problem is an issue and how it can be mitigated. Aggarwal et al. [ABL+18] conclude that the Hashcash proof of work used by Bitcoin is "relatively resistant to substantial speedup by quantum computers in the next 10 years". This is because, even given optimistic estimates about the near-term development of quantum computers, classical ASIC mining will continue to outperform quantum mining despite the quadratic speedup offered by Grover search. Nonetheless, in the medium term it remains possible that quantum computers will comprise a significant portion of mining power. Aggarwal et al. suggest a potential mitigation using an alternative proof of work called Momentum, which is based on finding collisions in a hash function. This is a more "quantum-resistant" proof of work, in the sense that there is a classical algorithm for finding collisions in a random function $\{0,1\}^m \to \{0,1\}^n$ in time $T = O(2^{n/2})$, whereas any quantum algorithm requires at least $\Omega(2^{n/3}) = \Omega(T^{2/3})$ queries, giving a somewhat smaller speedup than for the Hashcash proof of work.

However, this proposal has a significant drawback. The best algorithms for finding collisions [QD89] have the property that the probability of finding a collision increases quadratically in the running time of the algorithm. As a result, the Momentum proof of work suffers from the Superlinearity Problem in both the quantum and the classical settings.

Behnia et al. [BPOY21] proposes an alternative to Momentum based on the fact that the best known classical and quantum algorithms for the problem of finding a short vector in a lattice (in certain parameter regimes) have similar time complexities: 1.22^n and 1.20^n, respectively. Unlike for preimage and collision finding, the success probability of the "sieving" algorithms achieving these complexities has not been analysed as a function of running time. Nonetheless, it is easily seen that this function is at least quadratic.

Cojocaru et al. [CGK+20] give a formal asymptotic analysis of the security of the "Bitcoin backbone" protocol against quantum adversaries. They conclude that the protocol remains secure so long as malicious quantum parties control a very small fraction of the total computing power. We note that this fraction actually becomes *smaller* as the puzzle becomes more difficult, suggesting that an increase in *classical* computational power may help quantum attackers.

Prior work on the Superlinearity Problem. We are not aware of any prior work that considers the Quantum Superlinearity Problem. However, prior work has considered why superlinearity *in general* is a problem in blockchain protocols. Chen, Papadimitriou, and Roughgarden [CPR19] conduct a game-theoretic analysis of *allocation rules* in proof-of-work blockchains. An allocation rule is *proportional* if each miner (in expectation) receives reward proportional to their contribution relative to total network power. They show that the proportional allocation rule is the unique rule which satisfies both sybil- and collusion-resistance.

Nerem and Gaur [NG21] give another analysis of the impact of quantum mining on Bitcoin, examining the threshold at which quantum mining asympotically outperforms classical mining. Their analysis considers a simple Markov model of a single quantum miner Q competing with classical miners, and shows that when Q holds a very small fraction of total network power,

[3] E.g., verifiable random functions and verifiable delay functions.

Q only has a *linear* speedup over classical mining. They note that this analysis arises from an approximation that holds only if Q is weak compared to the network; that is, a quantum computer with a constant fraction of network power can still get superlinear rewards.

1.1 Technical overview

Next, we briefly summarise the techniques underlying our results.

Impossibility result. First, recall the Hashcash proof of work, and why it is subject to superlinear quantum attack. A Hashcash challenge consists of a hash function $h\colon \{0,1\}^\lambda \to \{0,1\}^\lambda$ and a target set $S \subseteq \{0,1\}^\lambda$. A proof π is some x such that $h(x) \in S$. If h is pseudorandom, then by evaluating h at t random points, we find such an x with probability roughly $t|S|/2^\lambda$. Moreover, no classical algorithm does better than this. But there is a *quantum* algorithm (Grover search) that makes t calls to h and finds such an x with probability $\Omega(t^2|S|/2^\lambda)$.

In general, a proof-of-work scheme may not have this form; indeed, the alternative constructions described above are quite different. Our key observation is that the Grover attack applies to all *proportional* proofs of work; i.e., where the probability of producing a valid proof scales *linearly* with the amount of work.

More precisely, let $\mathsf{Work}(c,t;r)$ denote the "honest" proof of work algorithm, for challenge c with time bound t using randomness r, and let $\mathsf{Verify}(c,\pi)$ be the verification algorithm. By proportionality, $\Pr_r[\mathsf{Verify}(c,\mathsf{Work}(c,t;r))] = \Theta(t)$.[4] To construct a quantum attack, we define a function $f_c(r) = \mathsf{Verify}(c,\mathsf{Work}(c,t_0;r))$, where t_0 is the smallest time such that Work outputs a proof with positive probability p. Then by running t iterations of Grover search on f_c we obtain with probability $\Omega(t^2)$ a string r such that $\mathsf{Work}(c,t_0;r)$ is a valid proof of work.

New proof-of-work construction. We present a proof-of-work construction in the random beacon model which avoids both the Quantum Advantage and Superlinearity Problems. Similarly to Hashcash, we compute $h(x)$ for many random x. We store each pair $(x,h(x))$ in a data structure sorted by $h(x)$. When the beacon value $\beta \in \{0,\dots,2^\lambda-1\}$ arrives, we search through the data structure for x such that $|\beta - h(x)|$ is minimized, and publish x as the proof. The verification algorithm accepts if $|\beta - h(x)|$ is below some specified difficulty threshold.

Intuitively, this circumvents the impossibility because proofs cannot be verified until after the beacon value arrives. We show this formally in Theorem 4.7 by modelling h as a quantum-accessible random oracle, using Zhandry's compressed oracle technique. In other words, we show that the honest classical mining algorithm is asymptotically optimal for both classical and quantum miners.

1.2 Summary of contributions

1. We identify and initiate the study of the Quantum Superlinearity Problem.
2. *Impossibility.* We prove an impossibility, namely, that the Superlinearity Problem is inherent in a large class of proofs of work, encompassing all existing definitions and constructions to our knowledge (Section 2).
3. *Possibilities.* We analyse our impossibility theorem to systematically highlight new proof-of-work approaches and other alternatives that may avoid the Superlinearity Problem, as open directions for future work (Section 3).

[4] This is not strictly true: the left hand side is bounded by 1 whereas t grows without bound. A refined definition of proportionality (Theorem 2.4) handles this issue.

4. *Construction.* We offer a new proof-of-work construction in a random-beacon model that provably avoids the Superlinearity Problem, and discuss remaining challenges of integrating it into a consensus protocol (Section 4).

2 The Quantum Superlinearity Problem is inherent

In this section, we show that the Superlinearity Problem is inherent in a broad class of proofs of work. Section 2.1 introduces the necessary definitions, and Section 2.2 presents the impossibility theorem and proof.

2.1 A broad definition of proofs of work

Our aim in this context is not to propose a canonical definition of proof of work that is somehow better than the scattered existing definitions in the literature, but rather, to be as *as inclusive as possible* — since the broader the definition, the stronger the impossibility.

Relation to existing proof-of-work definitions. Our definition generalizes many existing definitions of proofs of work from the literature, including Dwork and Naor's seminal "pricing functions" [DN92], Chen et al.'s "client puzzles" [CMSW09], Miller et al.'s "scratch-off puzzles" [MKKS15], Garay et al.'s "signatures of work" [GKP20], and Ball et al's proofs of work [BRSV18].

Our definition is also compatible with Jakobsson and Juels' proofs of work [JJ99], which additionally discusses *interactive* proofs of work. Our definition is incomparable with Stebila et al.'s "client puzzles" [SKR+11], which have a non-public verification algorithm. Section 3 provides further discussion of these and other model variants not captured by our definition.

Finally, our definition captures all proof-of-work constructions for blockchains of which we are aware, including Hashcash [Bac02] and Equihash [BK17].

Definition 2.1. A *proof of work* is parametrized by a *proof space* $\Pi = \{\Pi_\lambda\}_{\lambda\in\mathbb{N}}$ and a *challenge space* $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda\in\mathbb{N}}$, and consists of a triple of algorithms $\mathsf{POW} = (\mathsf{Gen}, \mathsf{Work}, \mathsf{Verify})$ with the following syntax. Gen and Work may be randomized; Verify is deterministic. Gen and Verify must be efficient (i.e., polynomial time).

- $\mathsf{Gen}(1^\lambda, \gamma)$ takes as input a security parameter λ (in unary) and difficulty parameter $\gamma \in [0, 1]$ and outputs a *challenge* $c \in \mathcal{C}_\lambda$.
- $\mathsf{Work}(c, t)$ takes as input a challenge $c \in \mathcal{C}_\lambda$ and *time parameter* $t \in \mathbb{N}$, and outputs a proof of work $\pi \in \Pi_\lambda$. The time complexity of Work is $t \cdot W_\gamma(\lambda)$, where W_γ is a polynomial.
- $\mathsf{Verify}(c, \pi)$ takes as input a challenge $c \in \mathcal{C}_\lambda$ and a candidate proof of work $\pi \in \Pi_\lambda$ and outputs $b \in \{0, 1\}$.

Remark 1. We can omit λ and γ from the input to $\mathsf{Work}, \mathsf{Verify}$ since we can assume that c includes them both without loss of generality. We sometimes write $\mathsf{Work}(c, t; r)$ to explicitly denote the randomness r of the Work algorithm.

Remark 2. $W_\gamma(\lambda)$ may be thought to represent the minimum time required to produce a valid proof for difficulty parameter γ with positive probability.

Next, we define the *reward function* of a proof of work (Definition 2.2), which relates the likelihood of obtaining a valid proof to (honest) work done.

Definition 2.2. A proof of work POW has *reward function* ρ if the probability of generating a valid proof by running Work with time parameter t with respect to difficulty parameter γ is negligibly close to $\rho(\gamma, t)$ with overwhelming probability. That is, for any $\lambda, t \in \mathbb{N}$, $\gamma \in [0,1]$, there exists a negligible function ε such that

$$\Pr_{c \leftarrow \mathsf{Gen}(1^\lambda, \gamma)} \left[\left| \Pr_r \left[b = 1 \quad : \quad \begin{array}{l} \pi \leftarrow \mathsf{Work}(c, t; r) \\ b \leftarrow \mathsf{Verify}(c, \pi) \end{array} \right] - \rho(\gamma, t) \right| \geq \varepsilon(\lambda) \right] \leq \varepsilon(\lambda) \ . \tag{1}$$

The bound (1) states (in the contrapositive) that with all but negligible probability over challenges, the probability of obtaining a valid proof in time t is very close to $\rho(\gamma, t)$, where γ is the difficulty parameter.

Next, we define *smoothness* and *proportionality* of reward functions. Informally, a reward function is *smooth* if no matter how hard the difficulty is set ($\gamma \to 0$), there is a positive probability of obtaining a valid proof after one time-step of computation.[5] A reward function is *proportional* if for all small enough γ (i.e., for all hard enough difficulty settings), the probability of obtaining a valid proof scales approximately linearly with computation, up to a positive upper bound.[6] Note that proportionality implies smoothness.

Definition 2.3. For $\alpha, \beta \in (0,1]$, we say a reward function ρ is (α, β)-*smooth* if for any $\gamma \in [0, \beta]$, $\rho(\gamma, 1) \geq \alpha \cdot \gamma$.

Definition 2.4. We say a reward function ρ is *proportional* if there exist $\alpha, \beta \in (0,1]$ such that for any $\gamma \in [0, \beta]$, $\alpha \cdot \min(\gamma t, 1) \leq \rho(\gamma, t) \leq \gamma t$.

Remark 3. Hashcash-like "progress-free" [BK17] proofs of work have $\rho(\gamma, t) = 1 - (1-\gamma)^t$: a proportional reward structure according to our definition.

Finally, we define the *hardness* of a proof of work. Informally, POW is *classically (resp., quantumly)* (μ_C, ρ')-*hard* if any classical (resp. quantum) algorithm running in time μ_C computes a valid proof of work with probability ρ'. Definitions 2.5–2.7 state these properties formally.

Definition 2.5 (Classical hardness)**.** A proof of work POW is *classically* (μ_C, ρ')-*hard* if for any classical two-part adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that $\mathcal{A}_1$ runs in polynomial time and $\mathcal{A}_2$ runs in time at most $\mu_C(t)$, for any $t \in \mathbb{N}$, and $\gamma \in [0,1]$, there is a negligible function ε such that

$$\Pr \left[b = 1 \quad : \quad \begin{array}{l} z \leftarrow \mathcal{A}_1(1^\lambda, t, \gamma) \\ c \leftarrow \mathsf{Gen}(1^\lambda, \gamma) \\ \pi \leftarrow \mathcal{A}_2(z, c) \\ b \leftarrow \mathsf{Verify}(c, \pi) \end{array} \right] \leq \rho'(\gamma, t) + \varepsilon(\lambda) \ . \tag{2}$$

Definition 2.6 (Quantum hardness)**.** A proof of work POW is *quantumly* (μ_C, ρ')-*hard* if for any quantum adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that $\mathcal{A}_1$ runs in polynomial time and $\mathcal{A}_2$ runs in time at most $\mu_Q(t)$, for any $t, n \in \mathbb{N}$ and $\gamma \in [0,1]$, there is a negligible function ε such that (5) holds.

Definition 2.7 (Hardness)**.** A proof of work POW is (μ_C, μ_Q, ρ')-*hard* if it is classically (μ_C, ρ')-hard and quantumly (μ_Q, ρ')-hard.

[5] This rules out deterministic proofs of work (whose reward functions are 0-1).

[6] The upper bound is necessary since probabilities are upper-bounded by 1.

2.2 Impossibility result

Our impossibility result relates reward functions to achievable quantum speedup. Any proof of work with a smooth reward function must allow quantum adversaries a probability of obtaining a valid proof that is quadratic in work done (Theorem 2.8). It follows that any proof of work with a *proportional* reward function must admit a quadratic quantum speedup (Corollary 2.9).

Theorem 2.8. *If a proof of work with an (α, β)-smooth reward function is quantum (μ_Q, ρ')-hard then there exists $\kappa \in [0,1]$ such that for all $\gamma \in [0, \beta]$ and all sufficiently large $t \in \mathbb{N}$, $\rho'(\gamma, t) \geq \kappa \cdot \min(\mu_Q(t)^2 \alpha\gamma, 1)$.*

Proof. Let $\mathsf{POW} = (\mathsf{Gen}, \mathsf{Work}, \mathsf{Verify})$ be a proof of work with reward structure ρ. A quantum attack on POW proceeds as follows. Consider the algorithm $f_c(r) = \mathsf{Verify}(c, \mathsf{Work}(c, 1; r))$. A single invocation of f_c runs in time $\mathsf{poly}(\lambda)$. By Eq. (1), $\Pr_r[f_c(r) = 1] \geq p(\gamma, 1) - \mathsf{negl}(\lambda)$; denote this probability by p_0.

We use Grover search to find r such that $f_c(r) = 1$, stopping at time $\mu_Q(t)$. This algorithm makes $\mu_Q(t)/\mathsf{poly}(\lambda)$ queries to the f_c-oracle. Denote the probability that this algorithm finds a winning choice of r by $p(\lambda, \gamma, t)$. By the standard analysis of Grover search, we have that $p(\lambda, \gamma, t) \geq \kappa(\lambda) \cdot \min(\mu_Q(t)^2 p_0, 1)$ for some $\kappa(\lambda) \colon \mathbb{N} \to [0,1]$ and all sufficiently large t.

Suppose that POW is (μ_Q, ρ')-hard; then there exists λ_0 such that for all $\lambda \geq \lambda_0$ and all γ, t it holds that $p(\lambda, \gamma, t) \leq \rho'(\gamma, t)$. Hence in particular for all sufficiently large t it holds that $\rho'(\gamma, t) \geq \kappa(\lambda_0) \cdot \min(\mu_Q(t)^2 p_0, 1)$. Noting that $p_0 \geq p(\gamma, 1)/2$ for large enough λ yields the theorem. □

Corollary 2.9. *If a proof of work with proportional reward function ρ is quantum (μ_Q, ρ)-hard then $\mu_Q = O(\sqrt{t})$.*

Proof. Let POW be a proof of work with proportional reward function ρ that is quantum (μ_Q, ρ)-hard. Let α, β be as in Theorem 2.4. By Theorem 2.8, there exist $\kappa \in [0,1], t_0 \in \mathbb{N}$ such that for all $\gamma \in [0, \beta]$ and $t \geq t_0$, $\gamma t \geq \rho(\gamma, t) \geq \kappa \cdot \min(\mu_Q(t)^2 \cdot \alpha \cdot \gamma, 1)$. Setting $\gamma = 1/\mu_Q(t)^2$ we see that there exists t_0' such that for all $t \geq t_0'$, $t \geq \kappa \cdot \mu_Q(t)^2$. □

The results in this section are already broad enough to encompass all prior proof-of-work constructions of which we are aware. Yet an even broader impossibility result can be shown using essentially the same techniques: namely, covering a larger class of proofs of work in which: (1) the scheme may depend on an additional setup function, (2) Gen and Work may take arbitrary auxiliary input, and (3) all algorithms have access to an arbitrary quantum-accessible oracle.[7] We present the narrower impossibility here for simplicity of exposition; we offer formal definitions and a theorem for the broader impossibility in Appendix A.

3 Towards Bypassing the Superlinearity Problem

In this section, we analyze the scope of our impossibility to identify potential paths forward. Informally restated, Theorem 2.8 tells us that the Superlinearity Problem is inherent in any proof of work that has *all* of the following properties:

1. the prover takes some input,
2. then performs classical computational work for a time t

[7] A *quantum-accessible* oracle permits queries in superposition. Formally, for a classical oracle f, we allow access to the unitary mapping $|x, y\rangle \mapsto |x, f(x) \oplus y\rangle$.

3. to output a proof π of polynomial size,
4. which verifies successfully with a probability (over Work) increasing in t,
5. using a verification algorithm that is deterministic, runs in polynomial time, and takes (only) c and π as input.
6. in the standard model or a quantum-accessible-oracle-based model.

Towards bypassing the Superlinearity Problem, then, we consider how proofs of work could be designed *not* to satisfy any of the properties listed above — since such proof-of-work schemes would fall outside the scope of Theorem 2.8.

Let us consider each of the listed properties in turn. Property 1 — namely, that the prover takes some input — admits no meaningful modification.

Property 2 is arguably inherent to the concept of a proof of (classical computational) work; as such, it highlights the possibilities of avoiding the Superlinearity Problem by turning to proofs of resources other than classical computational work. One option is to consider proofs of *quantum* computational work; this would preclude classical mining, of course, but might be acceptable in scenarios where efficient quantum computers are sufficiently widely available. More broadly, as also noted in Section 1, proof-of-work alternatives are an already thriving research area for which our results offer novel additional motivation.

Property 3 states that there is a proof string of polynomial size, raising the possibilities of having an interactive proof or having a proof string of superpolynomial size. Interactive proofs are not suitable for existing blockchain-based consensus systems, and would often incur prohibitive communication overhead; that said, interactive proofs of work could still be an interesting direction for future work.[8] Superpolynomial size, however, is unacceptable for efficiency.

Property 4 states the chance of obtaining a valid proof increases with work, a condition that seems inherent to the notion of a proof of work in the blockchain context. That is, for incentive-compatible mining in blockchain systems, the probability of obtaining a valid proof must increase with work. Property 4 raises the possibility of a proof of work where the probability of obtaining a valid proof still increases with work, but does not depend only on the random coins of Work. (For example, it might also depend on a miner's private information or on the randomness of an oracle.)

Property 5 seems arguably necessary in the blockchain context as existing blockchain networks rely crucially on verification being efficiently publicly computable and agreed on by everyone. Still, alternative models where Verify is not publicly computable (say, because it is keyed) may be worth exploring.[9]

Property 6 means that oracle-based cryptographic models such as the random oracle model and the common reference string (CRS) model will not take us outside the scope of Theorem 2.8's impossibility. However, other common non-standard cryptographic models could — such as a timed random beacon model or assuming a sybil-free public-key infrastructure (PKI) — as could oracle-based models that are not quantum-accessible. The latter would include, for example, oracles implemented under certain trusted hardware models, or oracles implemented by third parties (or networks of parties).

In summary, we have highlighted the following preliminary avenues for exploration towards designing a proof of work that falls outside the scope of our impossibility result (and thus may not suffer from the Superlinearity Problem).

[8] Jakobsson and Juels proposed a definition that includes interactive protocols [JJ99].

[9] Stebila et al. proposed a definition where verification is keyed [SKR+11].

A. Proofs of quantum work

B. Proof-of-work alternatives

C. Interactive proofs of work

D. Probability of obtaining a valid proof does not depend only on Work

E. Proofs of work with non-public verification

F. Non-standard non-oracle-based cryptographic models (e.g., beacon; PKI)

G. Oracle-based cryptographic models that are not quantum-accessible

In Section 4, we elaborate a preliminary proposal based on F.

4 A new proof of work in a random beacon model

We provide a proof-of-work construction in a timed random beacon model that provably avoids the Superlinearity Problem. Our protocol relies on the existence of a "beacon" that outputs a random string at regular time intervals. We prove our protocol's security in the quantum random oracle model (QROM). Then, in Sections 4.1 and 4.2, we highlight the significant challenges that seem inherent in integrating our proof of work into a realistic blockchain.

Model We define variants in the random beacon model of *proof of work* (Definition 4.1)[10], *reward structure* (Definition 4.2)[11], and (ρ, μ)-*hardness* (Definition 4.3)[12].

Definition 4.1 (Proof of work with beacon)**.** A *proof of work* is parametrized by a *proof space* Π and a *challenge space* $\mathcal{C}$, and consists of a triple of algorithms $\mathsf{POW} = (\mathsf{Gen}, \mathsf{Work}, \mathsf{Verify})$ with the following syntax. All the algorithms may be randomized. Gen and Verify must be efficient; Work need not be.

- $\mathsf{Gen}(1^\lambda, \gamma)$ takes as input a security parameter λ (in unary) and public parameters $\gamma \in [0,1]$ and outputs a *challenge* $c \in \mathcal{C}$.
- $\mathsf{Work}(c, t)$ takes as input a challenge $c \in \mathcal{C}$ and *time parameter* $t \in \mathbb{N}$, runs for time at most $t \cdot W_\gamma(\lambda)$ for some polynomial W_γ, and outputs a private state $\mathcal{D} \in \{0,1\}^*$.
- $\mathsf{Choose}^{\mathcal{D}}(\beta)$ takes as input a state $\mathcal{D}$ (as an oracle) and an auxiliary input $\beta \in \{0,1\}^{\mathsf{poly}(\lambda)}$ and outputs a proof of work $\pi \in \Pi$.
- $\mathsf{Verify}(c, \pi, \beta)$ takes as input a challenge $c \in \mathcal{C}$, a candidate proof of work $\pi \in \Pi$, beacon input $\beta \in \{0,1\}^{\mathsf{poly}(\lambda)}$ and outputs $b \in \{0,1\}$.
- $\mathsf{Beacon}(1^\lambda)$ takes as input security parameter λ (in unary) and outputs a beacon value $\beta \in \{0,1\}^{\mathsf{poly}(\lambda)}$.

Definition 4.2 (Reward function with beacon)**.** A proof of work and allocation algorithm POW has *reward function* ρ if for any $t \in \mathbb{N}$ and $\gamma \in [0,1]$, there is a negligible function ε such that

$$\Pr_{c \leftarrow \mathsf{Gen}(1^\lambda, \gamma)} \left[\left| \Pr \left[b = 1 \quad : \begin{array}{l} \mathcal{D} \leftarrow \mathsf{Work}(c, t) \\ \beta \leftarrow \mathsf{Beacon}(1^\lambda) \\ \pi \leftarrow \mathsf{Choose}^{\mathcal{D}}(\beta) \\ b \leftarrow \mathsf{Verify}(c, \pi, \beta) \end{array} \right] - \rho(\gamma, t) \right| > \varepsilon(\lambda) \right] < \varepsilon(\lambda) \ . \tag{3}$$

[10] See Definition 2.1 for standard-model definition.
[11] See Definition 2.2 for standard-model definition.
[12] See Definition 2.7 for standard-model definition

Definition 4.3 (Hardness with beacon)**.** A proof of work with beacon POW is classically $(\mu_{C,2}, \mu_{C,3}, \rho')$-hard if for any classical three-part adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ such that $\mathcal{A}_1$ runs in polynomial time, $\mathcal{A}_2$ runs in time at most $\mu_{C,2}(t) \cdot W(\lambda)$ and $\mathcal{A}_3$ runs in time at most $\mu_{C,3}(t) \cdot W(\lambda)$, and $\gamma \in [0,1]$, there is a negligible function ε such that

$$\Pr\left[\begin{array}{lll} & & z_1 \leftarrow \mathcal{A}_1(1^\lambda, \gamma) \\ & & c \leftarrow \mathsf{Gen}(1^\lambda, \gamma) \\ b = 1 & : & z_2 \leftarrow \mathcal{A}_2(z_1, c) \\ & & \beta \leftarrow \mathsf{Beacon}(1^\lambda) \\ & & \pi \leftarrow \mathcal{A}_3^{z_2}(c, \beta) \\ & & b \leftarrow \mathsf{Verify}(c, \pi, \beta) \end{array}\right] < \rho'(\gamma, t) + \varepsilon(\lambda) \ . \tag{4}$$

We define quantum $(\mu_{Q,2}, \mu_{Q,3}, \rho')$-hardness analogously.

Next we give our construction of a proof of work in the random beacon model, and then prove that it satisfies Theorems 4.2 and 4.3.

Construction 4.4. Challenge space $\mathcal{C} = \{0,1\}^\lambda \times \mathbb{N}$; proof space $\Pi = \{0,1\}^{2\lambda}$.

- $\mathsf{Gen}(1^\lambda, \gamma)$ samples uniform $\alpha \leftarrow \{0,1\}^\lambda$, computes $\tau = \gamma \cdot 2^{\lambda-1}$, and outputs $c = (1^\lambda, \tau, \alpha)$.
- $\mathsf{Work}(c, t)$ repeats the following t times: choose a random nonce $r \in \{0,1\}^\lambda$, compute $y = h_\alpha(r)$, and store (r, y) in a table $\mathcal{D}$ sorted by y (interpreted as an integer). It then outputs $\mathcal{D}$.
- $\mathsf{Choose}^{\mathcal{D}}(\beta)$ outputs $\pi = \arg\min_{(r,y)\in\mathcal{D}} |y - \beta|$ (interpreting y, β as integers.
- $\mathsf{Verify}(c, (r, y), \beta)$ accepts if $y = h_\alpha(r)$ and $|y - \beta| < \tau$.
- $\mathsf{Beacon}(1^\lambda)$ outputs uniformly random $\beta \in \{0,1\}^\lambda$.

Lemma 4.5. *When h is pseudorandom and runs in time $t_h(\lambda)$, Theorem 4.4 is a proof of work with reward structure $\rho(\gamma, t) = 1 - (1-\gamma)^t$, which achieves proportional representation.*

Proof. We show that replacing h_α with a random function h yields reward structure ρ. For a random h, the probability that for a random r and any β that $|h(r) - \beta| < \tau$ is γ. The probability that at least one proof succeeds out of t is then $1 - (1-\gamma)^t$, since these are independent events.

We now show that ρ achieves proportional representation. For the upper bound, observe that for all γ, t, $\rho(\gamma, t) \le \gamma t$. For the lower bound, suppose first that $\gamma \ge 1/t$; then $(1-\gamma)^t \le (1 - 1/t)^t \le 1/e$. Now suppose instead that $\gamma t < 1$; then $1 - (1-\gamma)^t \ge \gamma t - \frac{1}{2}(\gamma t)^2 \ge \gamma t/2$. □

QROM preliminaries. We introduce the technical background for our security proof in the quantum random oracle model (QROM) [BDF+11]. We omit standard quantum information definitions (e.g., states, unitaries) (see [NC16, §I.2]).

Let $\mathcal{A}$ be an algorithm that makes t quantum queries to an oracle $h\colon X \to \{0,1\}^\lambda$ and outputs a pair $(x, y) \in X \times \{0,1\}^\lambda$. Then there exist unitary transformations $U_1, \ldots, U_t$ and a quantum state $|\psi_0\rangle$ such that for any h, x, y,

$$\Pr\left[(x, y) \leftarrow \mathcal{A}^h\right] = \left\| \langle x, y | \, U_t O_h U_{t-1} O_h \cdots U_1 O_h \, |\psi_0\rangle \right\|^2 \ ,$$

where O_h is the unitary with action $O_h |x, y\rangle = |x, y \oplus h(x)\rangle$ for all x, y and $\oplus$ is the bitwise XOR.

Compressed oracle technique. We make use of Zhandry's compressed oracle technique [Zha19]. Let $X \rightharpoonup \{0,1\}^\lambda$ be the set of partial functions from X to $\{0,1\}^\lambda$. For $D\colon X \rightharpoonup \{0,1\}^\lambda$, let $\mathsf{supp}(D)$ be the set of $x \in X$ for which $D(x)$ is defined. Let $\mathcal{D}$ be a quantum register supported on states $|D\rangle$ for $D\colon X \rightharpoonup \{0,1\}^\lambda$. The key lemma of the compressed oracle technique follows.

Lemma 4.6 ([Zha19])**.** *There exists a unitary $\mathcal{O}$ such that for all $R \subseteq X \times \{0,1\}^\lambda$, letting $\Pi_R = \sum_{D,\exists(x,y)\in R, D(x)=y} |D\rangle\langle D|_{\mathcal{D}}$, $\mathcal{A}^{\mathcal{O}} = U_t\mathcal{O}\cdots U_1\mathcal{O}$:*

$$\Pr\left[\begin{array}{c|c} (x,y)\in R & h\leftarrow(X\to\{0,1\}^\lambda) \\ \wedge\, h(x)=y & (x,y)\leftarrow\mathcal{A}^h \end{array}\right] \leq \left\|\Pi_R\mathcal{A}^{\mathcal{O}}|\psi_0\rangle|\perp\rangle_{\mathcal{D}}\right\|^2 + O(2^{-\lambda}) .$$

Moreover, let $p_R = \max_x \Pr_y[(x,y)\in R]$. Then $\|[\Pi_R,\mathcal{O}]\| = O(\sqrt{p_R})$.

Theorem 4.7. *In the (quantum) random oracle model, Theorem 4.4 is (ρ,μ_C,μ_Q)-hard for $\mu_{C,2},\mu_{C,3} = \Theta(t)$, $\mu_{Q,2} = \Theta(t), \mu_{Q,3} = \Theta(\sqrt{t})$.*

Proof. The classical hardness proof is straightforward, and so we omit it.

Let $\mathcal{A} = (\mathcal{A}_1,\mathcal{A}_2,\mathcal{A}_3)$ be quantum oracle algorithms making at most t_1,t_2,t_3 queries respectively. Let $R(\alpha,\beta) = \{((\alpha,x),y) : |y-\beta| \leq \tau\}$; note that for all α,β, $p_{R(\alpha,\beta)} = p_R = 2\tau/2^\lambda$. By Theorem 4.6 the probability that $\mathcal{A}$ produces a valid proof of work is at most

$$\delta = \mathsf{E}_{\alpha,\beta}\|\Pi_{R(\alpha,\beta)}\mathcal{A}_3^{\mathcal{O}}(c,B)\mathcal{A}_2^{\mathcal{O}}(c)\mathcal{A}_1^{\mathcal{O}}|\psi_0\rangle|\perp\rangle\|^2 + O(2^{-\lambda}) .$$

Let $S(\alpha) = \{((\alpha,x),y) : x,y\in\{0,1\}^\lambda\}$ and $\bar{\Pi} = I - \Pi$. Then

$$\begin{aligned}
&\|\Pi_{R(\alpha,\beta)}\mathcal{A}_3^{\mathcal{O}}(\alpha,\beta)\mathcal{A}_2^{\mathcal{O}}(\alpha)\mathcal{A}_1^{\mathcal{O}}|\psi_0\rangle|\perp\rangle\| \\
&\quad\leq \|\Pi_{R(\alpha,\beta)}\mathcal{A}_3^{\mathcal{O}}(\alpha,\beta)\bar{\Pi}_{R(\alpha,\beta)}\mathcal{A}_2^{\mathcal{O}}(\alpha)\mathcal{A}_1^{\mathcal{O}}|\psi_0\rangle|\perp\rangle\| + \\
&\qquad \|\Pi_{R(\alpha,\beta)}\mathcal{A}_2^{\mathcal{O}}(\alpha)\bar{\Pi}_{S(\alpha)}\mathcal{A}_1^{\mathcal{O}}|\psi_0\rangle|\perp\rangle\| + \|\Pi_{S(\alpha)}\mathcal{A}_1^{\mathcal{O}}|\psi_0\rangle|\perp\rangle\|
\end{aligned}$$

by the triangle inequality. Now we bound each term in turn. For all α,β,

$$\begin{aligned}
&\|\Pi_{R(\alpha,\beta)}\mathcal{A}_3^{\mathcal{O}}(\alpha,\beta)\bar{\Pi}_{R(\alpha,\beta)}\mathcal{A}_2^{\mathcal{O}}(\alpha)\mathcal{A}_1^{\mathcal{O}}|\psi_0\rangle|\perp\rangle\| \\
&\leq \|\Pi_{R(\alpha,\beta)}\mathcal{A}_3^{\mathcal{O}}(\alpha,\beta)\bar{\Pi}_{R(\alpha,\beta)}\| \\
&\leq \|\mathcal{A}_3^{\mathcal{O}}(\alpha,\beta)\Pi_{R(\alpha,\beta)}\bar{\Pi}_{R(\alpha,\beta)}\| + t_3\cdot\|[\Pi_R,\mathcal{O}]\| = O(t_3\sqrt{p_R}),
\end{aligned}$$

where the final equality follows by Lemma 4.6 and because $\Pi_{R(\alpha,\beta)}\bar{\Pi}_{R(\alpha,\beta)} = 0$.

For the second term, observe that for any α and any state $|\varphi\rangle$ in the image of $I - \Pi_{S(\alpha)}$, the support of $\mathcal{D}$ in $\mathcal{A}_2^{\mathcal{O}}(\alpha)|\varphi\rangle$ is contained in the set

$$S = \{D : |\mathsf{supp}(D)\cap\{\alpha\|x : x\in\{0,1\}^\lambda\}| \leq t_2\} .$$

Hence if we measure $\mathcal{D}$, we obtain $D\in S$ with probability 1. For all such D, $\Pr_\beta[\exists x, D(\alpha\|x) = \beta] \leq t_2\cdot p_R$. Hence for all α,

$$\mathsf{E}_\beta\|\Pi_{R(\alpha,\beta)}\mathcal{A}_2^{\mathcal{O}}(\alpha)\bar{\Pi}_{S(\alpha)}\mathcal{A}_1^{\mathcal{O}}|\psi_0\rangle|\perp\rangle\|^2 \leq t_2\cdot p_R .$$

A similar argument shows that $\mathsf{E}_\alpha\|\Pi_{S(\alpha)}\mathcal{A}_1^{\mathcal{O}}|\psi_0\rangle|\perp\rangle\|^2 = O(t_1/2^\lambda)$. Then

$$\delta = O((t_1+\tau t_2+\tau t_3^2)/2^\lambda) = O((t_2+t_3^2)\cdot\frac{\tau}{2^\lambda}) + \mathsf{negl}(\lambda) .$$

□

4.1 Challenges of protocol integration

Given our new proof-of-work construction, one might hope to "plug it in" to a Bitcoin-like protocol and thus resolve the Quantum Superlinearity Problem. Unfortunately, integrating our proof of work into a decentralised consensus protocol seems to present non-trivial further challenges. Next, we briefly elaborate on these, guided by sketches of simple but natural failed attempts.

Why not, for instance, rely on the beacon to keep time (say, one block per beacon output) and have miners publish proofs of work after each beacon value? A fundamental issue with this approach is takeover attacks: a malicious miner could create an alternate history on a fork or an entire alternate chain knowing the beacon values after the fact, and obtain a chain indistinguishable from — or of higher quality than — the honestly derived chain for any network participant who is newly joining or joining after an offline period. This problem seems difficult to mitigate when network participants are not almost always online.

Inspired by this observation, we might propose a variant protocol that requires miners to publish commitments to their candidate proofs in each time-step, and only considers valid those proofs that miners can prove were committed "on-chain". To achieve this, individual miners must be able to decommit their own proofs (in a publicly verifiable way). This constraint appears to preclude the natural approach of committing to all miners' commitments with a single Merkle root. But then, storing commitment information on-chain that scales with the total number of miners would incur an impractical bandwidth cost.

When the validity of a proof of work is effectively dependent on when it was computed, and participants are not always online, it is arguably inherent that a consensus protocol dependent on such a proof of work must record some timing information. The problem we have highlighted lies in recording timing information *even for unsuccessful attempts* at block mining, which creates impractical bandwidth demands. We would be interested to see future work exploring new approaches to integrating timing-dependent proofs of work into blockchains.

4.2 Challenges of beacon implementation

Other challenges of using our proof-of-work construction to avoid the Superlinearity Problem in practice arise from the fact that it relies on a random beacon.

For our application, the beacon must be post-quantum secure. Unfortunately, current plausible approaches to implementing a random beacon are pre-quantum: for example, one common approach is to use verifiable delay functions [BBBF18], efficient known constructions of which are pre-quantum [Pie19b, Wes20]. Post-quantum secure variants exist but are less efficient and less well understood [Kho20]. Furthermore, many VDF-based beacon applications require basically that nobody can compute the beacon value *before* a given time (e.g., lotteries). A consensus protocol built on a proof-of-work like ours likely requires the stronger guarantee that everyone learns the value at roughly the same time. While this problem could be mitigated by having honest parties who learn the beacon value verifiably broadcast it, complications arise in the presence of network delays or powerful adversaries that can withhold the beacon.

5 Conclusion

We have identified the Quantum Superlinearity Problem in post-quantum proof-of-work blockchains and proven that it is inherent in a large class of proofs of work encompassing existing approaches. By analyzing our impossibility result, we have suggested a range of approaches or alternatives to proofs of work that may have the potential to avoid the Superlinearity Problem. We have explored one such approach in more detail, proposing a new proof-of-work construction in a

random-beacon model that provably avoids the Superlinearity Problem; and we provide discussion of the significant challenges that seem to remain in integrating our new proof of work into a realistic consensus protocol.

Finally, we have highlighted several open problems and directions for future research to improve our understanding of post-quantum blockchains in light of the Superlinearity Problem, as follows.

1. Explore new models of proofs of work that avoid the Superlinearity Problem (including, but not limited to, the directions A–G noted in Section 3).

2. Explore how such new proofs of work (including, but not limited to, our construction in Section 4) can be integrated into decentralised consensus protocols that avoid the Superlinearity Problem.

3. Explore post-quantum implementations of a public random beacon.

Acknowledgments

We are grateful to Thaddeus Dryja for the conversation that sparked this research, and to Chris Peikert for a helpful discussion at early stages of the work.

SP's work on this project was supported by a 2021–22 Computing Innovation Fellowship, funded by the National Science Foundation under Grant #2127309 to the Computing Research Association, by Cornell Tech's Digital Life Initiative, and by the MIT Media Lab's Digital Currency Initiative.

References

[AAB+19] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), October 2019.

[ABL+18] Divesh Aggarwal, Gavin Brennen, Troy Lee, Miklos Santha, and Marco Tomamichel. Quantum attacks on bitcoin, and how to protect against them. *Ledger*, 3, Oct. 2018.

[AW22] Nick Arnosti and S. Matthew Weinberg. Bitcoin: A natural oligopoly. *Manag. Sci.*, 68(7):4755–4771, 2022.

[Bac02] Adam Back. Hashcash - a denial of service counter-measure, 2002. http://www.hashcash.org/papers/hashcash.pdf.

[BBBF18] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 757–788. Springer, 2018.

[BCEM15] Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–38, May 2015.

[BDE+22] Maxime Buser, Rafael Dowsley, Muhammed F. Esgin, Shabnam Kasra Kermanshahi, Veronika Kuchta, Joseph K. Liu, Raphaël C.-W. Phan, and Zhenfei Zhang. Post-quantum verifiable random function from symmetric primitives in pos blockchain. In Vijayalakshmi Atluri, Roberto Di Pietro, Christian Damsgaard Jensen, and Weizhi Meng, editors, *Computer Security - ESORICS 2022 - 27th European Symposium on Research in Computer Security, Copenhagen, Denmark, September 26-30, 2022, Proceedings, Part I*, volume 13554 of *Lecture Notes in Computer Science*, pages 25–45. Springer, 2022.

[BDF+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Proceedings of the 17th International Conference on the Theory and Application of Cryptology and Information Security*, ASIACRYPT '11, pages 41–69, 2011.

[BK17] Alex Biryukov and Dmitry Khovratovich. Equihash: Asymmetric proof-of-work based on the generalized birthday problem. *Ledger*, 2:1–30, Apr. 2017.

[Bli18] Marc Blinder. Making cryptocurrency more environmentally sustainable. Harvard Business Review (Online), 2018. https://hbr.org/2018/11/making-cryptocurrency-more-environmentally-sustainable.

[BMC+15] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. Sok: Research perspectives and challenges for bitcoin and cryptocurrencies. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 104–121. IEEE Computer Society, 2015.

[BMW22] Roger E. Barton, Christopher J. McNamara, and Michael C. Ward. Are cryptocurrencies securities? the SEC is answering the question. Reuters, June 2022. https://www.reuters.com/legal/transactional/are-cryptocurrencies-securities-sec-is-answering-question-2022-03-21 [https://perma.cc/32DQ-PB4J].

[Bol20] Andreas Bolfing. Post-Quantum Blockchains. In *Cryptographic Primitives in Blockchain Technology: A mathematical introduction.* Oxford University Press, 09 2020.

[BPOY21] Rouzbeh Behnia, Eamonn W. Postlethwaite, Muslum Ozgur Ozmen, and Attila Altay Yavuz. Lattice-based proof-of-work for post-quantum blockchains. In

DPM/CBT@ESORICS, volume 13140 of *Lecture Notes in Computer Science*, pages 310–318. Springer, 2021.

[BRSV18] Marshall Ball, Alon Rosen, Manuel Sabin, and Prashant Nalini Vasudevan. Proofs of work from worst-case assumptions. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 789–819. Springer, 2018.

[Cas22] Amy Castor. Why ethereum is switching to proof of stake and how it will work, 2022. https://www.technologyreview.com/2022/03/04/1046636/ethereum-blockchain-proof-of-stake [https://perma.cc/U957-V7X7].

[CGK+20] Alexandru Cojocaru, Juan Garay, Aggelos Kiayias, Fang Song, and Petros Wallden. Post-quantum blockchain proofs of work, 2020.

[CMSW09] Liqun Chen, Paul Morrissey, Nigel P. Smart, and Bogdan Warinschi. Security notions and generic constructions for client puzzles. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*, pages 505–523. Springer, 2009.

[Coi22] CoinMarketCap. Today's cryptocurrency prices by market cap, June 2022. https://coinmarketcap.com [https://perma.cc/9ARA-AXBQ].

[Cou] Council of the European Union. Digital finance: agreement reached on european crypto-assets regulation (mica). Press release. https://www.consilium.europa.eu/en/press/press-releases/2022/06/30/digital-finance-agreement-reached-on-european-crypto-assets-regulation-mica [https://perma.cc/36NR-DQVQ].

[CP19] Bram Cohen and Krzysztof Pietrzak. The chia network blockchain, 2019. https://www.chia.net/wp-content/uploads/2022/07/ChiaGreenPaper.pdf.

[CPR19] Xi Chen, Christos H. Papadimitriou, and Tim Roughgarden. An axiomatic approach to block rewards. In *AFT*, pages 124–131. ACM, 2019.

[DN92] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In Ernest F. Brickell, editor, *Advances in Cryptology - CRYPTO '92, 12th Annual International Cryptology Conference, Santa Barbara, California, USA, August 16-20, 1992, Proceedings*, volume 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer, 1992.

[ES18] Ittay Eyal and Emin Gün Sirer. Majority is not enough: bitcoin mining is vulnerable. *Commun. ACM*, 61(7):95–102, 2018.

[Eth] Ethereum.org. Proof-of-stake (pos). https://ethereum.org/en/developers/docs/consensus-mechanisms/pos [https://perma.cc/FB7M-SZU2].

[FF20] Tiago M. Fernández-Caramés and Paula Fraga-Lamas. Towards post-quantum blockchain: A review on blockchain cryptography resistant to quantum computing attacks. *IEEE Access*, 8:21091–21116, 2020.

[GBE+18] Adem Efe Gencer, Soumya Basu, Ittay Eyal, Robbert van Renesse, and Emin Gün Sirer. Decentralization in bitcoin and ethereum networks. In Sarah Meiklejohn and Kazue Sako, editors, *Financial Cryptography and Data Security - 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26 - March 2, 2018, Revised Selected Papers*, volume 10957 of *Lecture Notes in Computer Science*, pages 439–457. Springer, 2018.

[GKP20] Juan A. Garay, Aggelos Kiayias, and Giorgos Panagiotakos. Consensus from signatures of work. In Stanislaw Jarecki, editor, *Topics in Cryptology - CT-RSA 2020 - The Cryptographers' Track at the RSA Conference 2020, San Francisco, CA, USA, February 24-28, 2020, Proceedings*, volume 12006 of *Lecture Notes in Computer Science*, pages 319–344. Springer, 2020.

[GYB22] Erin Griffith and David Yaffe-Bellany. Bitcoin plummets below $20,000 for first time since late 2020. New York Times, June 2022. https://www.nytimes.com/2022/06/18/technology/bitcoin-20000.html.

[JJ99] Markus Jakobsson and Ari Juels. Proofs of work and bread pudding protocols. In Bart Preneel, editor, *Secure Information Networks: Communications and Multimedia Security, IFIP TC6/TC11 Joint Working Conference on Communications and Multimedia Security (CMS '99), September 20-21, 1999, Leuven, Belgium*, volume 152 of *IFIP Conference Proceedings*, pages 258–272. Kluwer, 1999.

[Kho20] Dmitry Khovratovich. Stark-based vdfs, 2020. https://ethresear.ch/t/stark-based-vdfs/8052 [https://perma.cc/6RYN-R28J].

[Kle22] Zoe Kleinman. Bitcoin: Why is the largest cryptocurrency crashing? BBC, June 2022. https://www.bbc.co.uk/news/technology-61796155 [https://perma.cc/6PNV-9AZ7].

[KN12] Sunny King and Scott Nadal. PPCoin: Peer-to-peer crypto-currency with proof-of-stake, 2012.

[KO19] Sinan Küfeoğlu and Mahmut Özkuran. Bitcoin mining: A global review of energy and power demand. *Energy Research & Social Science*, 58:101273, 2019.

[LaM21] Brian LaMacchia. The long road ahead to transition to post-quantum cryptography. *Commun. ACM*, 65(1):28–30, December 2021.

[LBS22] Sishan Long, Soumya Basu, and Emin Gün Sirer. Measuring miner decentralization in proof-of-work blockchains. *CoRR*, abs/2203.16058, 2022.

[MKKS15] Andrew Miller, Ahmed E. Kosba, Jonathan Katz, and Elaine Shi. Nonoutsourceable scratch-off puzzles to discourage bitcoin mining coalitions. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, pages 680–691. ACM, 2015.

[MP21] Michele Mosca and Marco Piani. 2021 quantum threat timeline report, 2021. https://globalriskinstitute.org/publications/2021-quantum-threat-timeline-report [https://perma.cc/8AU5-2JDC].

[Nak09] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. May 2009. http://www.bitcoin.org/bitcoin.pdf.

[Nat22] National Institute of Standards and Technology (NIST). Post-quantum cryptography, 2022. https://csrc.nist.gov/projects/post-quantum-cryptography [https://perma.cc/6U4S-VEDW].

[NC16] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information (10th Anniversary edition)*. Cambridge University Press, 2016.

[NG21] Robert R. Nerem and Daya R. Gaur. Conditions for advantageous quantum bitcoin mining, 2021.

[Nxt14] Nxt Community. Nxt whitepaper, 2014. https://www.jelurida.com/sites/default/files/NxtWhitepa

[Osb22] Margaret Osborne. Bitcoin could rival beef or crude oil in environmental impact. Smithsonian Magazine, 2022. https://www.smithsonianmag.com/smart-news/bitcoin-could-rival-beef-or-crude-oil-in-environmental-impact-180980877 [https://perma.cc/8WJH-NVPU].

[Pie19a] Krzysztof Pietrzak. Proofs of catalytic space. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 59:1–59:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[Pie19b] Krzysztof Pietrzak. Simple verifiable delay functions. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 60:1–60:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[PKF+18] Sunoo Park, Albert Kwon, Georg Fuchsbauer, Peter Gazi, Joël Alwen, and Krzysztof Pietrzak. Spacemint: A cryptocurrency based on proofs of space. In Sarah Meiklejohn and Kazue Sako, editors, *Financial Cryptography and Data Security - 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26 - March 2, 2018, Revised Selected Papers*, volume 10957 of *Lecture Notes in Computer Science*, pages 480–499. Springer, 2018.

[Pro17] Protocol Labs. Filecoin: A decentralized storage network, 2017. https://filecoin.io/filecoin.pdf.

[QD89] Jean-Jacques Quisquater and Jean-Paul Delescaille. How easy is collision search. new results and applications to DES. In *CRYPTO*, volume 435 of *Lecture Notes in Computer Science*, pages 408–413. Springer, 1989.

[Shi21] Shinobi. How centralized is bitcoin mining really? Bitcoin Magazine, dec 2021. https://bitcoinmagazine.com/business/is-bitcoin-mining-centralized.

[SKR+11] Douglas Stebila, Lakshmi Kuppusamy, Jothi Rangasamy, Colin Boyd, and Juan Manuel González Nieto. Stronger difficulty notions for client puzzles and denial-of-service-resistant protocols. In Aggelos Kiayias, editor, *Topics in Cryptology - CT-RSA 2011 - The Cryptographers' Track at the RSA Conference 2011, San Francisco, CA, USA, February 14-18, 2011. Proceedings*, volume 6558 of *Lecture Notes in Computer Science*, pages 284–301. Springer, 2011.

[SSP13] Elaine Shi, Emil Stefanov, and Charalampos Papamanthou. Practical dynamic proofs of retrievability. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung,

editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 325–336. ACM, 2013.

[Vra17] Harald Vranken. Sustainability of bitcoin and blockchains. *Current Opinion in Environmental Sustainability*, 28:1–9, 2017. Sustainability governance.

[Wag] Jeremiah Wagstaff. Subspace: A solution to the farmer's dilemma. https://drive.google.com/file/d/1v847u_XeVf0SBz7Y7LEMXi72QfqirstL/view [https://perma.cc/W33J-CQNK].

[Wes20] Benjamin Wesolowski. Efficient verifiable delay functions. *J. Cryptol.*, 33(4):2113–2147, 2020.

[Whi22] White House Office of Science and Technology Policy (OSTP). Climate and energy implications of crypto-assets in the united states. 2022. https://www.whitehouse.gov/wp-content/uploads/2022/09/09-2022-Crypto-Assets-and-Climate-Report.pdf [https://perma.cc/7DDQ-KYX9].

[Zha19] Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In *CRYPTO (2)*, volume 11693 of *Lecture Notes in Computer Science*, pages 239–268. Springer, 2019.

A Generalized impossibility result

Definition A.1. An *auxiliary-input proof of work with optional setup* is parametrized by a *proof space* $\Pi = \{\Pi_\lambda\}_{\lambda\in\mathbb{N}}$, a *challenge space* $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda\in\mathbb{N}}$, and an efficiently samplable auxiliary distribution, $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ and consists of a quadruple of algorithms $\mathsf{POW} = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Work}, \mathsf{Verify})$ with the following syntax. Setup, Gen, and Work may be randomized; Verify is deterministic. Setup, Gen, and Verify must be efficient.

- $\mathsf{Setup}(1^\lambda, a_0)$ takes as input a security parameter λ and auxiliary input $a_0 \in \{0,1\}^{\mathsf{poly}(\lambda)}$, and outputs parameters pp.
- $\mathsf{Gen}(\mathsf{pp}, \gamma, a_1)$ takes as input the parameters pp, difficulty parameter $\gamma \in [0,1]$, and auxiliary input $a_1 \in \{0,1\}^{\mathsf{poly}(\lambda)}$, and outputs *challenge* $c \in \mathcal{C}_\lambda$.[13]
- $\mathsf{Work}(c, t, a_2)$ takes as input a challenge $c \in \mathcal{C}_\lambda$, *time parameter* $t \in \mathbb{N}$, and auxiliary input $a_2 \in \{0,1\}^{\mathsf{poly}(\lambda)}$, and outputs a proof of work $\pi \in \Pi_\lambda$. The time complexity of Work is $t \cdot W_\gamma(\lambda)$, where W_γ is a polynomial.
- $\mathsf{Verify}(c, \pi)$ takes as input a challenge $c \in \mathcal{C}_\lambda$ and a candidate proof of work $\pi \in \Pi_\lambda$, and outputs $b \in \{0,1\}$.

Definition A.2 (Reward function with auxiliary input and setup)**.** An auxiliary-input proof of work with optional setup, POW, has *reward function* ρ if the probability of generating a valid proof by running Work with time parameter t with respect to difficulty parameter γ is negligibly close to $\rho(\gamma, t)$ with overwhelming probability. That is, for any $\lambda, t \in \mathbb{N}$, $\gamma \in [0,1]$, there exists a negligible function ε such that

$$\Pr_{\substack{(a_0,a_1,a_2)\leftarrow\mathcal{X}\\ \mathsf{pp}\leftarrow\mathsf{Setup}(1^\lambda,a_0)\\ c\leftarrow\mathsf{Gen}(\mathsf{pp},\gamma,a_1)}}\left[\left|\Pr_r\left[\begin{array}{lll} b=1 & : & \pi\leftarrow\mathsf{Work}(c,t,a_2;r)\\ & & b\leftarrow\mathsf{Verify}(c,\pi)\end{array}\right]-\rho(\gamma,t)\right|\geq\varepsilon(\lambda)\right]\leq\varepsilon(\lambda)\ .$$

As before, informally, POW is *classically* (μ_C, ρ')*-hard* if any classical algorithm running in time μ_C computes a valid proof of work with probability ρ', and POW is *quantumly* (μ_Q, ρ')*-hard* if the analogous statement holds for quantum algorithms.

Definition A.3 (Classical hardness with auxiliary input and setup)**.** An auxiliary-input proof of work with optional setup, POW, is *classically* (μ_C, ρ')*-hard* if for any classical two-part adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that $\mathcal{A}_1$ runs in polynomial time and $\mathcal{A}_2$ runs in time at most $\mu_C(t)$, for any $t \in \mathbb{N}$, $\gamma \in \{0,1\}^*$, and $a_0, a_1, a_2 \in \{0,1\}^{\mathsf{poly}(\lambda)}$, there is a negligible function ε such that

$$\Pr\left[\begin{array}{lll} b=1 & : & \mathsf{pp}\leftarrow\mathsf{Setup}(1^\lambda,a_0)\\ & & z\leftarrow\mathcal{A}_1(\mathsf{pp},t,\gamma,a_2)\\ & & c\leftarrow\mathsf{Gen}(1^\lambda,\gamma,a_1)\\ & & \pi\leftarrow\mathcal{A}_2(z,c)\\ & & b\leftarrow\mathsf{Verify}(c,\pi)\end{array}\right]\leq\rho'(\gamma,t)+\varepsilon(\lambda)\ . \tag{5}$$

Definition A.4 (Quantum hardness with auxiliary input and setup)**.** An auxiliary-input proof of work with optional setup, POW, is *quantumly* (μ_C, ρ')*-hard* if for any quantum adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that $\mathcal{A}_1$ runs in polynomial time and $\mathcal{A}_2$ runs in time at most $\mu_Q(t)$, for any $\in \mathbb{N}$, $\gamma \in \{0,1\}^*$, and $a_0, a_1, a_2 \in \{0,1\}^{\mathsf{poly}(\lambda)}$, there is a negligible function ε such that (5) holds.

[13] Without loss of generality, c may be considered to contain pp.

Definition A.5 (Hardness with auxiliary input and setup)**.** An auxiliary-input proof of work with optional setup, POW, is (μ_C, μ_Q, ρ')-*hard* if it is classically (μ_C, ρ')-hard and quantumly (μ_Q, ρ')-hard.

Theorem A.6. *If an auxiliary-input proof of work with optional setup with reward structure ρ is quantum (μ_Q, ρ')-hard then there exists $\kappa \in [0,1]$ such that for all $\gamma \in [0,1]$ and all sufficiently large $t \in \mathbb{N}$, $\rho'(\gamma, t) \geq \kappa \cdot \min(\mu_Q(t)^2 \rho(\gamma, 1), 1)$.*