Quantum Proofs of Space

Jack Maxfield, The Aerospace Corporation SURF Fellow

Professor Thomas Vidick

October 7, 2020

1 Abstract

Proofs of space have been developed as a more energy-efficient alternative to proofs of work protocols. They have applications in blockchain technology and deterrence of denial-of-service attacks. A proof of space protocol allows one to demonstrate they have dedicated a significant amount of disk space to a problem. We hope to leverage the properties of quantum information to create a quantum proof of space, a quantum analogue to proofs of space in which use of quantum memory is demonstrated. After developing a definition for quantum proofs of space, we developed a quantum proof of space protocol and proved its correctness and security.

2 Introduction

In cryptography, there is a notion of a “proof of work” (PoW), an interactive protocol between two parties which allows one party to demonstrate to the other that they have allocated a nontrivial amount of computational resources to a problem. These protocols have a number of uses, particularly blockchain applications and denial of service and spam email prevention. For example, imagine anyone who wishes to send an email must include a proof of work in the body of their message. This would not affect legitimate users who send only several emails per day, but might be a large burden to spammers wishing to send out millions of emails, thereby deterring spam.

There is a related notion called a “proof of space” (PoS) [2], which demonstrates use of storage rather than computation time. In a PoS, one party demonstrates they are allocating a nontrivial amount of storage space over time. As each proof of work is the result of a long, energy intensive computation, proofs of space are often thought of as a more environmentally friendly alternative to proofs of work. More formally, a proof of space is an interactive protocol between two parties, the prover and the verifier, in which the prover convinces the verifier she is using a large amount of storage.

In short, our goal was to create a quantum analogue to proofs of space, a protocol in which the prover demonstrates they are wasting a large amount of *quantum* storage. Like the classical protocol, we additional demand succinct communication between the prover and verifier, for example, requiring that the transmissions between the two parties are much smaller than the quantum storage used by the prover. We hope to leverage the properties of quantum information to create quantum PoS protocols with better security than is classically possible.

Conducting literature review, there is no existing work on quantum PoS. We developed a rigorous definition for a quantum PoS, including a security definition. Next, we presented examples of quantum PoS protocols and proved their security and correctness.

3 Quantum Proof of Space Definition

Our first goal was to rigorously pin down what a quantum PoS ought to look like. We came up with the following definition, analogous to the classical PoS presented in [2]: a quantum proof of space is a two-stage interactive protocol between two quantum Turing machines, the prover P and verifier V. The protocol consists of the following two stages.

3.1 Stages

Initialization is an interactive protocol in which both machines receive n (an efficiency parameter), γ (the security parameter), and potentially other parameters. In this stage, an honest prover would generate a large quantum state and store it until execution. Denoting the parameters $\mathsf{prm} = (n, \gamma, ...)$, execution results in $(\Phi, S) = \langle V, P\rangle(\mathsf{prm})$, where Φ, S are potentially quantum. Note that V can "abort" in this phase by outputting $\Phi =\perp$.

Execution is an interactive protocol in which P and V have access to the values from initialization. In this stage, an honest prover would demonstrate to the verifier that she has stored the values generated during initialization. Like the classical protocol, P has no output, while V can accept or reject.

$$(\{\mathsf{accept}, \mathsf{reject}\}, \varnothing) \leftarrow \langle \mathsf{V}(\Phi), \mathsf{P}(S) \rangle (\mathsf{prm})$$

Note that unlike the classical protocol [2], we do not demand that the execution phase can be repeated, since a single execution may alter the quantum state stored by either party.

3.2 Completeness, security, and efficiency

We model the dishonest prover $\tilde{\mathsf{P}}$ as follows. A N_0 prover $\tilde{\mathsf{P}}$'s quantum storage is bounded by N_0 after the initialization phase. We place no restriction on $\tilde{\mathsf{P}}$'s classical storage or classical computation abilities, except for those imposed by the efficiency requirement below. Note that N_0 is a function of n, and informally, the above protocol is an N_0 quantum proof of space if an honest prover is accepted and a dishonest N_0 prover $\tilde{\mathsf{P}}$ is rejected.

Completeness The honest prover is accepted with near certainty, i.e.

$$\Pr[\mathsf{out} = \mathsf{accept} \mid (\Phi, S) \leftarrow \langle \mathsf{V}, \mathsf{P} \rangle (\mathsf{prm}), (\mathsf{out}, \varnothing) \leftarrow \langle \mathsf{V}(\Phi), \mathsf{P}(S) \rangle (\mathsf{prm})] = 1 - \mathrm{negl}(\gamma)$$

where γ is the security parameter.

Soundness Any N_0 dishonest prover $\tilde{\mathsf{P}}$ is accepted with negligible probability.

$$\Pr[\mathsf{out} = \mathsf{accept} \mid (\Phi, S) \leftarrow \langle \mathsf{V}, \tilde{\mathsf{P}} \rangle (\mathsf{prm}), (\mathsf{out}, \varnothing) \leftarrow \langle \mathsf{V}(\Phi), \tilde{\mathsf{P}}(S) \rangle (\mathsf{prm})] = \mathrm{negl}(\gamma)$$

Efficiency The verifier V is efficient, taking time polynomial in n. The prover runs in time $\text{poly}(N_0)$ during initialization and execution. Informally, we would also like for N_0 to grow quickly in n, i.e. $N_0 = \Omega(n^2)$ or $N_0 = \Omega(2^n)$, but this is not required.

4 The Trivial Protocol

In this section, we briefly introduce the "trivial protocol" which is perhaps the simplest example of a quantum PoS. It is nearly identical to the "Dimension test protocol" from Chao and Reichardt [1], both of which are stated here, as well as a reduction from the Dimension Test to the Trivial Quantum PoS protocol considered in this paper.

Chao and Reichardt Dimension Test

Let $N \geq 1$ (the "storage bound") be an integer and $\alpha \in [0, 1/2)$.

Initialization The verifier chooses $B_0 \in \{0, 1\}$ and $S \in \{0, 1\}^N$, both uniformly at random. For $j = 1, \ldots, N$, he prepares and sends the prover a qubit in state $H^{B_0} |S_j\rangle$, where $H = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right)$ is the Hadamard matrix.

Execution The verifier announces his basis choice B to the prover. Based on B, the prover measures her system, and outputs a guess $S' \in \{0, 1\}^N$.

The prover passes the test if and only if S and S' match on at least $(1 - \alpha)N$ bits.

Trivial Quantum PoS Protocol

Let $N \geq 1$ (the "storage bound") be an integer and $\alpha \in [0, 1/2)$.

Initialization The verifier chooses $B \in \{0,1\}^N$ and $S \in \{0,1\}^N$, both uniformly at random. For $j = 1, \ldots, N$, he prepares and sends the prover a qubit in state $H^{B_j} |S_j\rangle$, where $H = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right)$ is the Hadamard matrix.

Execution The verifier announces his basis choice B to the prover. Based on B, the prover measures her system, and outputs a guess $S' \in \{0,1\}^N$.

The prover passes the test if and only if S and S' match on at least $(1-\alpha)N$ bits.

The difference between the Trivial Quantum PoS Protocol and the "Dimension test protocol" from [1] is that the above protocol a different basis for each qubit, while the "Dimension test protocol" chooses one basis at random and encodes all qubits in it. Here, we argue the Trivial Protocol is at least as hard to win as the Chao and Reichardt Dimension Test.

Theorem 4.1. *Consider a prover* P *which succeeds in the Trivial Protocol with probability* p*, i.e.*

$$\frac{1}{2^N} \sum_{B' \in \{0,1\}^N} \Pr[\textit{succeed}|B = B'] = p$$

Then there is a prover P' *which uses equal storage to* P *and succeeds in the Dimension Test with probability* $\geq p$*.*

Proof. For a string $B' \in \{0,1\}^n$, let $\overline{B}'$ be the bitwise inversion of B', i.e. if $B' = 10101$, then $\overline{B}' = 01010$. We can write P's probability of success as

$$\frac{1}{2^{N-1}} \sum_{B' \in 0\{0,1\}^{N-1}} \left[\frac{1}{2} \Pr[\text{succeed}|B = B'] + \frac{1}{2} \Pr[\text{succeed}|B = \overline{B}']\right] = p$$

where $0\{0,1\}^{N-1}$ is the set of bit-strings of length N that begin with 0. There must exist some $B' \in \{0,1\}^N$ such that

$$\frac{1}{2} \Pr[\text{succeed}|B = B'] + \frac{1}{2} \Pr[\text{succeed}|B = \overline{B}'] \geq p$$

Now, imagine a prover P' which, during the Dimension Tets, behaves nearly identically to P, except with the following changes:

During initialization When the ith qubit is received, P' applies $H^{B'_i}$ before processing it.

During execution If $B_0 = 0$, P' responds how P would if $B = B'$. If $B_0 = 1$, P' responds how P would if $B = \overline{B}'$.

Effectively, during Initialization, the state received by P' is coded in the bases B' if $B_0 = 0$, and $\overline{B}'$ if $B_0 = 1$. During execution, if $B_0 = 0$, the prover's chance of winning is exactly $\Pr[\text{succeed}|B = B']$, and likewise if $B_0 = 1$, the prover's chance of winning is $\Pr[\text{succeed}|B = \overline{B}']$ As $\frac{1}{2} \Pr[\text{succeed}|B = B'] + \frac{1}{2} \Pr[\text{succeed}|B = \overline{B}'] \geq p$, P' succeeds in the Dimension Test with probability $\geq p$ despite using no more quantum storage. □

After Step 1, the prover holds some classical-quantum state in $\mathcal{H}_C \otimes \mathcal{H}_Q$, where $\mathcal{H}_C$ is a finite dimensional Hilbert space containing the classical information and likewise $\mathcal{H}_Q$ contains the quantum part. Chao and Reichardt prove the following statement about the security of the Dimension Test [1].

Theorem 4.2 (Chao & Reichardt Dimensionality Lower Bound)**.** *If the prover passes with probability p, then*

$$\log \dim H_Q \geq N - 2H(p) - 2p \log F - 2(1-p) \log(2^N - F)$$

where $F = \sum_{i=1}^{\alpha N} \binom{N}{i}$ and $H(x) = -x \log x - (1-x) \log(1-x)$ is the binary entropy function.

Corollary 4.2.1. *Theorem 4.2 applies to the Trivial Protocol.*

Proof. Apply Theorem 4.1. □

The Trivial Protocol is secure, but it is quite inefficient, as all states must be generated by the verifier, sent to the prover, and measured by the prover. In the next section, we will explore a more efficient variant of this protocol in which measures for only a small subset of qubits are requested during execution.

5 Efficient Variants of the Trivial Protocol

In this section, we will analyze the following protocol. The following notation is used: $[N] \coloneqq \{1, 2, ..., N\}$. S is a string of bits, and we may index subsequences of S by writing e.g. S_A for $A \subseteq [N]$. For example, if $S = 10101$ and $A = \{1, 3, 5\}$, then $S_A = 111$.

Efficient Quantum PoS Protocol

Let $N \geq 1$ (the "storage bound") and $C \geq 1$ (the "challenge size") be integers.

Initialization The verifier chooses $B \in \{0, 1\}$ and $S \in \{0, 1\}^N$, both uniformly at random. For $j = 1, \ldots, N$, he prepares and sends the prover a qubit in state $H^B |S_j\rangle$, where $H = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right)$ is the Hadamard matrix.

Execution The verifier randomly chooses a subset $A \in \binom{[N]}{C}$ uniformly at random. The verifier announces his basis choice B to the prover. Based on B, the prover measures qubits with indices in A, and outputs a guess $S' \in \{0, 1\}^C$.

The prover passes the test if and only if $S_A = S'$.

Completeness and efficiency clearly hold for this protocol. After Initialization, the prover holds some classical-quantum state in $\mathcal{H}_C \otimes \mathcal{H}_Q$, with quantum part $\rho \in \mathcal{H}_Q$. For each $A \in \binom{[N]}{C}$, there is some corresponding POVM $\{M_{AT}\}$ (indexed by $T \in \{0, 1\}^C$) the prover makes on ρ to respond to the challenge. For any subset $A \in \binom{[N]}{C}$, we define p_{succ_A} as the prover's chance of passing when challenged on that subset, i.e. $p_{\text{succ}_A} = \operatorname{Tr} M_{AS_A}\rho$.

We will give two proofs of security of this protocol under two different sets of assumptions.

5.1 From Strong Assumptions about the Prover

First, we will make the following assumption: the prover handles all qubits she receives independently. Assume the prover's quantum storage is bounded by kN qubits for $k \in [0, 1]$, and that C is a non-constant (i.e. not bounded by any constant) function of N. The following is true:

Theorem 5.1. *Let p_{win} be the win probability of any prover with the above limitations. If $k \in [0, 1)$, then $p_{win} \to 0$ as $N \to \infty$. If $k = 1$, the prover can pass with probability 1.*

Proof. First, note that during Initialization, each time a qubit is received, the prover can either store it or immediately make some measurement and store classical information. If she stores the state, she can correctly answer the query with probability 1. If not, as the states $|0\rangle, |1\rangle, |+\rangle$, and $|-\rangle$ are not orthogonal, they cannot be distinguished perfectly, so she has a chance $\omega < 1$ of answering queries correctly (the exact value of ω is not relevant for proving the asymptotic bound).

We consider a variant of the above protocol where, when the verifier generates the random subset $A \subseteq [N]$, he picks elements of $[N]$ *with replacement* and sends only the unique picks to the prover, so the size of A may be less than C. Briefly, this game is strictly easier for the prover to pass than the above game, because she can easily expand any A of size $< C$ to one of size C, drawn from the same distribution as the original game, then employ her strategy from the original game and win with the same probability. However, this variant is easier to analyze asymptotically.

Consider the random variable Y, the number of qubits in the random sample which are not stored, and thus can each be answered with probability at most ω.

The prover's chance of winning is given by

$$p_{\text{win}} \leq \sum_{c=0}^{\infty} \Pr[Y = c]\omega^c.$$

To show $p_{\text{win}} \to 0$, since $\omega < 1$, we first argue that $\Pr[Y = c] \to 0$ for any fixed c.

$$\Pr[Y = c] \leq \binom{C}{c}(1-k)^c k^C \propto \binom{C}{c} k^C \to 0$$

as $\binom{C}{c}$ is a polynomial of degree c, while k^C shrinks exponentially. To show why this is sufficient, for any $\varepsilon > 0$, we can pick N_0 such that $\sum_{i=N_0}^{\infty} \omega^i < \varepsilon/2$, and by the above observation, $N_1 \geq N_0$ such that for all $C \geq N_1$, $\sum_{i=0}^{N_0-1} \Pr[Y = c] < \varepsilon/2$. Then for all N such that $C(N) \geq N_1$, we have

$$\begin{aligned}
p_{\text{win}} &\leq \sum_{c=0}^{\infty} \Pr[Y = c]\omega^c \\
&= \sum_{c=0}^{N_0-1} \Pr[Y = c]\omega^c + \sum_{N_0}^{\infty} \Pr[Y = c]\omega^c \\
&\leq \sum_{c=0}^{N_0-1} \Pr[Y = c] + \sum_{N_0}^{\infty} \omega^c \\
&\leq \varepsilon/2 + \varepsilon/2 \\
&= \varepsilon
\end{aligned}$$

so $\lim_{N\to\infty} p_{\text{win}} = 0$.

Finally, if $k = 1$, the prover can store all states and measure at the end to pass with probability 1. □

5.2 From No Assumptions about the Prover

In this section, we prove security of the above protocol without assuming the prover handles all qubits independently. This will be done using a reduction to [1]. Specifically, we will show that if the prover can correctly answer any query of C qubits with probability p, then the prover can, by making multiple measurements, correctly answer for many qubits in the original protocol.

Suppose the prover is challenged on the values of some subset A of size C of qubits, i.e. $A \subset [N]$, $|A| = C$. Now, we assume the prover wins with probability p, so

$$\mathbb{E}_A\big[\operatorname{Tr} M_{AS_A}\rho\big] = p$$

First, we will argue there exists a family of subsets $B_1, ..., B_m$ such that $|\bigcup_i B_i|$ is large, and $\operatorname{Tr} M_{B_i S_{B_i}}\rho$ is large for all $1 \leq i \leq m$. This is shown via the following probabalistic argument. Fix some integer $n \geq 1$, and suppose we choose, uniformly at random and with replacement, n subsets $A_1, ..., A_n \subseteq [N]$, each of size C.

Lemma 5.2. *Fix $\varepsilon > 0$. With probability strictly greater than $1/2$, at least half of $A_1, ..., A_n$ have $p_{succ_{A_i}} \geq 1 - (2+\varepsilon)(1-p)$.*

Proof. Applying Markov's inequality to the random variable $1 - p_{\text{succ}_A}$, we get that

$$\Pr\big[1 - p_{\text{succ}_A} \geq (2+\varepsilon)(1-p)\big] \leq \frac{1-p}{(2+\varepsilon)(1-p)} = \frac{1}{2+\varepsilon} < \frac{1}{2}.$$

Thus, in any uniform random sample of subsets of $[n]$, it is more likely than not that at least half have $p_{\text{succ}} \geq 1 - (2+\varepsilon)(1-p)$. □

Lemma 5.3. *With probability at least $1/2$,*

$$\Big|\bigcup_i A_i\Big| \geq E - \sqrt{V}$$

where

$$E \coloneqq N\Big(1 - \Big[1 - \frac{C}{N}\Big]^n\Big) \tag{5.2.1}$$

and

$$V \coloneqq (N-E) + N(N-1)P - (N-E)^2 \tag{5.2.2}$$

where

$$P \coloneqq \left[\frac{(N-C)(N-C-1)}{N(N-1)}\right]^n.$$

Proof. Consider the random variable $Y \coloneqq |\bigcup_i A_i|$. It's enough to show the median of Y is at least $E - \sqrt{V}$. Using the fact that the mean and median differ by at most the standard deviation for a distribution with finite variance, we will show $\mathbb{E}[Y] = E$ and $\mathrm{var}[Y] = V$.

For the mean:

$$\mathbb{E}[Y] = \sum_{j=1}^{N} \mathbb{E}[\mathbb{1}_{j \in \bigcup_i A_i}]$$

For each j, the probability j does not appear in randomly chosen A is $1 - \frac{C}{N}$. The probability j appears in no A_i is thus $(1 - \frac{C}{N})^n$, so $\mathbb{E}[\mathbb{1}_{j \in \bigcup_i A_i}] = 1 - (1 - \frac{C}{N})^n$. Summing over j, we obtain the desired result.

For the variance: it is easier to analyze the random variable $Z \coloneqq N - Y$, the size of the subset not covered by A_is. Clearly $\mathrm{var}[Y] = \mathrm{var}[Z]$. First, we compute the probability P that for any $j \neq k$, that both j, k are not included in $\bigcup_i A_i$. For each A_i, this is $\binom{N-2}{C} / \binom{N}{C}$, so by independence

$$P \coloneqq \left[\frac{\binom{N-2}{C}}{\binom{N}{C}} \right]^n = \left[\frac{(N-C)(N-C-1)}{N(N-1)} \right]^n$$

Now, we can compute

$$\begin{aligned}
\mathrm{var}[Z] &= \mathbb{E}[Z^2] - \mathbb{E}[Z]^2 \\
&= \sum_{j=1}^{N} \sum_{k=1}^{N} \mathbb{E}[\mathbb{1}_{j,k \notin \bigcup_i A_i}] - (N-E)^2 \\
&= \sum_{j=1}^{N} \mathbb{E}[\mathbb{1}_{j \notin \bigcup_i A_i}] + \sum_{\substack{1 \leq j,k \leq N \\ j \neq k}} \mathbb{E}[\mathbb{1}_{j,k \notin \bigcup_i A_i}] - (N-E)^2 \\
&= (N-E) + N(N-1)P - (N-E)^2
\end{aligned}$$

completing the proof of Claim 2.

□

Now we apply the probabilistic argument.

Lemma 5.4. *Fix $\varepsilon > 0$. There exists a set $\{B_1, ..., B_m\}$ of elements of $\binom{[N]}{C}$ such that*

1. *For all i, $p_{succ_{B_i}} \geq 1 - (2+\varepsilon)(1-p)$.*
2. *$\{B_i\}$ have large intersection, i.e.*

$$\left| \bigcup_i B_i \right| \geq E - \sqrt{V} - \frac{nC}{2}$$

where E, V are as defined in Equations 5.2.1, 5.2.2 *respectively.*

3. $m \geq n/2$.

Proof. Note that the sum of probabilities of the events in Lemmas 5.2, 5.3 is strictly greater than 1, so the intersection of these events is non-empty, giving us a set $\{A_1, ..., A_n\}$ such that $|\bigcup_i A_i| \geq E - \sqrt{V}$ and at least half of A_i have $p_{\text{succ}_{A_i}} \geq 1 - (2+\varepsilon)(1-p)$. Discard A_i with $p_{\text{succ}_{A_i}} < 1 - (2+\varepsilon)(1-p)$ to obtain a subset $\{B_1, ..., B_m\} \subseteq \{A_1, ..., A_n\}$ where all $p_{\text{succ}_{B_i}} \geq 1 - (2+\varepsilon)(1-p)$ and $m \geq n/2$, and note that

$$\left|\bigcup_i B_i\right| \geq E - \sqrt{V} - \frac{nC}{2}$$

since at most $(n-m)C$ elements were discarded. □

Now, suppose the prover measures each of these states. Each measurement succeeds with probability at least $p' = 1-(2+\varepsilon)(1-p)$. We aim to bound the probability all measurements succeed. To do so, we use the *Gentle Measurement Lemma*, given as e.g. Lemma 9.4.1 in Wilde [3], stated here for convenience.

Lemma 5.5 (Gentle Measurement Lemma)**.** *Consider a density operator ρ and a measurement operator Λ where $0 \leq \Lambda \leq I$. The measurement operator could be an element of a POVM. Suppose that the measurement operator Λ has a high probability of detecting state ρ*

$$\operatorname{Tr} \Lambda\rho \geq 1 - \varepsilon,$$

where $\varepsilon \in [0,1]$ (the probability of detection is high if ε is close to zero). Then the post-measurement state

$$\rho' \equiv \frac{\sqrt{\lambda}\rho\sqrt{\lambda}}{\operatorname{Tr} \Lambda\rho}$$

is $2\sqrt{\varepsilon}$-close to the original state rho in trace distance:

$$|\rho - \rho'|_1 \leq 2\sqrt{\varepsilon}.$$

Thus, the measurement does not disturb the state ρ by much if ε is small.

We can view each measurement made by the prover as having two outcomes: "correct" and "incorrect", with "correct" having probability at least p' by construction of B_i. However, making these measurements on each B_i perturbs the original state by some amount so that subsequent measurements may not be as likely to proceed. We will let the prover make the measurements $\{M_{B_1}\}, \{M_{B_2}\}, \{M_{B_3}\}, ...$, and let ρ_i be the held quantum state after $B_1, ..., B_{i-1}$ have been measured, where $\rho_1 = \rho$.

Lemma 5.6. *Suppose that each measurements succeeds with probability at least p'. For all $1 \leq i \leq m$,*

$$|\rho_i - \rho_1|_1 \leq 2(i-1)\sqrt{1-p'}$$

so p_{all}, the probability of all measurements succeeding, is bounded below by

$$p_{all} \geq (p' - 2(n/2-1)\sqrt{1-p'})^{n/2}.$$

Proof. Induction on i and the Gentle Measurement Lemma. The base case, $i = 1$, is clear. The notation $M_{B_i}(\rho_j)$ will mean the state which exists after measurement $\{M_{B_i}\}$ is made on ρ_j.

$$\begin{aligned}
|\rho_i - \rho_1|_1 &= |M_{B_i}(\rho_{i-1}) - \rho_1|_1 \\
&= |M_{B_i}(\rho_{i-1}) - M_{B_i}(\rho_1) + M_{B_i}(\rho_1) - \rho_1|_1 \\
&\leq |M_{B_i}(\rho_{i-1}) - M_{B_i}(\rho_1)|_1 + |M_{B_i}(\rho_1) - \rho_1|_1 \text{ (triangle inequality)} \\
&\leq |\rho_{i-1} - \rho_1|_1 + |M_{B_i}(\rho_1) - \rho_1|_1 \text{ (contractivity of trace distance)} \\
&\leq 2(i-2)\sqrt{1-p'} + 2\sqrt{1-p'} \\
&= 2(i-1)\sqrt{1-p'}.
\end{aligned}$$

Thus, p_{all}, the probability of all measurements succeeding, is given by

$$\begin{aligned}
p_{\text{all}} &\geq \prod_{i=1}^{m} \left(p' - 2(i-1)\sqrt{1-p'}\right) \\
&\geq (p' - 2(m-1)\sqrt{1-p'})^m \\
&\geq (p' - 2(n/2-1)\sqrt{1-p'})^{n/2}
\end{aligned}$$

□

Theorem 5.7. *Suppose we run the "Efficient Quantum PoS Protocol" with some $N \geq C \geq 1$, and the prover can answer correctly with probability p, i.e.*

$$\mathbb{E}_{A \in \binom{[N]}{C}}[p_{succ_A}] = p.$$

Then for any $\varepsilon > 0$, $0 \leq \gamma < 1/2$, and $n \in \mathbb{Z}^+$, the prover's quantum storage $\log \dim \mathcal{H}$ *is bounded by*

$$\log \dim \mathcal{H}_\rho \geq N - 2H(p_{all}p_{greater}) - 2p_{all}p_{greater} \log F - 2(1 - p_{all}p_{greater}) \log(2^N - F)$$

where

$$M = E - \sqrt{V} - \frac{nC}{2}$$

for E, V defined in Equations 5.2.1, 5.2.2, *respectively, and*

$$p' := 1 - (2 + \varepsilon)(1 - p)$$

$$p_{all} := (p' - 2(n/2 - 1)\sqrt{1 - p'})^{n/2}$$

$$p_{greater} = \frac{1}{2^N} \sum_{j=\lceil \gamma(N-M) \rceil}^{(N-M)} \binom{N}{j}$$

$$\alpha := 1 - \frac{\gamma(N - M) + M}{N}$$

$$F := \sum_{i=1}^{\alpha N} \binom{N}{i}$$

Proof. We begin by applying Lemma 5.4 to obtain a set of measurable subsets $B_1, ..., B_m$ where

$$\left| \bigcup_i B_i \right| \geq E - \sqrt{V} - \frac{nC}{2} = M$$

and $p_{\text{succ}_{B_i}} \geq p'$ for all i. By measuring $B_1, ..., B_m$ in any order, by Lemma 5.6 we obtain correct measurements on at least M qubits with probability p_{all} as defined above. For the remaining $N - M$ qubits, guessing randomly yields at least $\lceil \gamma(N - M) \rceil$ correct measurements with probability p_{greater} as defined above. Thus, the probability of getting a fraction at least α correct is bounded below by $p_{\text{all}} p_{\text{greater}}$. Plugging into the bound from Theorem 2.1 from [1], we complete the reduction and obtain the desired bound. □

6 Further Work

The "efficient" protocol presented above has a number of undesirable properties. For one, all states must be generated by the verifier and sent to the prover, which is inefficient, as a good quantum PoS protocol ought to involve minimal quantum work on the part of the verifier. This section will describe a protocol based on quantum lightning, which is defined as follows [4].

For a given security parameter λ, a quantum lightning protocol is a means of sampling a pair $(\mathsf{Storm}, \mathsf{Ver})$, two polynomial time quantum algorithms. When run, Storm samples a "bolt" $|\psi\rangle$, a quantum state, while Ver is an algorithm which takes in a quantum state $|\varphi\rangle$ and checks if it is a bolt sampled using Storm. If not, it returns $\perp$, and if so, Storm returns a serial number s.

Security ("uniqueness") for quantum lightning is stated in terms of the following game between a challenger and any bolt generation procedure H [4]:

Quantum Lightning Security Game

1. The challenger samples a storm and verifier $(\mathsf{Storm}, \mathsf{Ver})$, and sends $(\mathsf{Storm}, \mathsf{Ver})$ to H.
2. H generates two quantum states $|B_1\rangle, |B_2\rangle$ and sends them to the challenger.
3. The challenger verifies each state with Ver and accepts if both states are accepted as bolts with equal serial numbers.

The quantum lightning protocol is said to be secure iff all H can win with only probability negligible in λ, the security parameter. Informally, it's difficult to create two bolts with the same serial number. This motivates the following protocol:

Quantum PoS from Quantum Lightning

Fix N, the storage bound, and λ, the security parameter.

Initialization The verifier samples a storm and verifier $(\mathsf{Storm}, \mathsf{Ver})$ and sends them to the prover. Using Storm and Ver, the prover generates N bolts $|B_1\rangle, ..., |B_N\rangle$, measures their serial numbers $s_1, ..., s_N$, and commits to them using a Merkle tree.

Execution The verifier chooses a bolt index uniformly at random and has the prover send over that bolt. He verifies it with Ver and accepts if its serial number is consistent with the Merkle tree commited earlier, which requires a small amount of classical communication between the prover and verifier.

This protocol has a number of desirable qualities. For one, the verifier is entirely classical during Initialization, and the amount of classical information exchanged is logarithmic in the storage bound, while the amount of quantum information exchanged depends only on λ. Further work on quantum PoS would focus on proving security of the above protocol, which initially appears difficult. For example, it's difficult to limit what a malicious prover might do with her states, and it is hard to rule out the prover compressing all of her states to a small quantum state. Without making assumptions such as those of Theorem 5.1, it will take much more work to prove (or disprove) security of this protocol.

References

[1] Rui Chao and Ben W Reichardt. Quantum dimension test using the uncertainty principle. *arXiv preprint arXiv:2002.12432*, 2020.

[2] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. In *Annual Cryptology Conference*, pages 585–605. Springer, 2015.

[3] Mark M Wilde. From classical to quantum shannon theory. *arXiv preprint arXiv:1106.1445*, 2011.

[4] Mark Zhandry. Quantum lightning never strikes the same state twice. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 408–438. Springer, 2019.