Proofs of Space with Maximal Hardness

Leonid Reyzin
Boston University*
https://www.cs.bu.edu/fac/reyzin

September 20, 2024

Abstract

In a proof of space, a prover performs a complex computation with a large output. A verifier periodically checks that the prover still holds the output. The security goal for a proof of space construction is to ensure that a prover who erases even a portion of the output has to redo a large portion of the complex computation in order to satisfy the verifier.

In existing constructions of proofs of space, the computation that a cheating prover is forced to redo is a small fraction (vanishing or small constant) of the original complex computation. The only exception is a construction of Pietrzak (ITCS 2019) that requires extremely depth-robust graphs, which result in impractically high complexity of the initialization process.

We present the first proof of space of reasonable complexity that ensures that the prover has to redo almost the entire computation (fraction arbitrarily close to 1) when trying to save even an arbitrarily small constant fraction of the space. Our construction is a generalization of an existing construction called SDR (Fisch, Eurocrypt 2019) deployed on the Filecoin blockchain. Our improvements, while general, also demonstrate that the already deployed construction has considerably better security than previously shown.

Technically, our construction can be viewed as amplifying predecessor-robust graphs. These are directed acyclic graphs in which every set of $\pi \cdot n$ nodes contains a subset of $\alpha_\pi \cdot n$ nodes whose induced subgraph has just one sink. We take a predecessor-robust graph with constant parameters (π, α_π), and build a bigger predecessor-robust graph with a near-optimal set of parameters and additional guarantees on sink placement, while increasing the degree only by a small additive constant.

1 Introduction

In a proof of persistent space [DFKP15], a verifier V wants to be convinced that a prover P is continuously using a lot of storage. To initialize a proof of space[1] instance,

*Work done, in part, while visiting Universitat Pompeu Fabra and Protocol Labs.

[1]We will omit the term "persistent" from now on; proofs of transient, as opposed to persistent, space were introduced in [ABFG14].

P takes a small instance identifier x, and generates a very large output $y = f(x)$. P also provides to V a commitment to y and a proof that the computation of f was correct or at least close to correct. V and P then periodically carry out a protocol called "execution": V queries random portions of y, which P returns together with the proof of their correctness.[2] To be confident that P is really using the storage, we need the following property: when storing less than all of y, it should be difficult for P to come up with portions of y in response to the queries of V.

What does "storing less than all of y" mean? Naturally, the prover is not limited to storing bits of y, and can store other values—for example, some intermediate values in the computation of $y = f(x)$. If P stores just a little bit less than $|y|$, perhaps answering V's queries is not so hard. A proof of space is thus characterized by a *space gap* $\varepsilon_{\mathsf{space}}$: if a cheating prover stores fewer than $(1 - \varepsilon_{\mathsf{space}}) \cdot |y|$ bits, then answering queries becomes "difficult"; but above $(1 - \varepsilon_{\mathsf{space}}) \cdot |y|$ storage, the proof of space provides no guarantees.

What does "difficult" mean? Answering the queries of V is never harder than simply recomputing $y = f(x)$ from the short input x. We will use the term *hardness gap* $\varepsilon_{\mathsf{hardness}}$ to denote the difference between the difficulty of initial computation of f and the difficulty of cheating. In other words, a proof of space should ensure that a prover who uses less than $(1 - \varepsilon_{\mathsf{space}}) \cdot |y|$ space will have computational cost at least $(1 - \varepsilon_{\mathsf{hardness}}) \cdot \mathrm{cost}(f(x))$ during execution. (Following the original definition of [DFKP15], our computational abstraction is a strong version of the random oracle model, and cost is measured as the total number of oracle queries.)

An optimal proof of space would get both the hardness gap $\varepsilon_{\mathsf{hardness}}$ and the space gap $\varepsilon_{\mathsf{space}}$ as close to 0 as possible, in order to make it maximally expensive for the adversary to even minimally reduce storage. A smaller $\varepsilon_{\mathsf{space}}$ ensures that the claimed storage is really being used, which is particularly important when power in a consensus protocol is allocated based on proofs of space [PPK+15, Pro17, CP19]. A smaller $\varepsilon_{\mathsf{hardness}}$ allows for increased time T_V between executions, by the following reasoning. In order to pass execution, a prover can choose to expend computational effort $(1 - \varepsilon_{\mathsf{hardness}}) \cdot \mathrm{cost}(f(x))$ or to store at least $(1 - \varepsilon_{\mathsf{space}}) \cdot |y|$ bits for time T_V; and a greater $\varepsilon_{\mathsf{hardness}}$ means that the tradeoff favors storage for a greater T_V. A greater T_V translates into reduced costs for the honest provers and verifiers, because they don't have to run execution as often.[3]

1.1 State of the Art

In most existing constructions, $\varepsilon_{\mathsf{space}}$ and $\varepsilon_{\mathsf{hardness}}$ are close to 1, meaning that a cheating prover can avoid most of the storage and most of the work. In such a case, it is more intuitive to talk about *space ratio* $r_{\mathsf{space}} = 1 - \varepsilon_{\mathsf{space}}$ and *hardness ratio* $r_{\mathsf{hardness}} = 1 - \varepsilon_{\mathsf{hardness}}$, because the guarantee provided by such a construction is that a prover who uses less than r_{space} fraction of the storage of $|y|$ will have to incur more than r_{hardness} fraction of

[2] There is only one construction [AAC+17] that manages to avoid proofs of correctness in both initialization and execution, and thus allows for initialization that is noninteractive, as well as very short messages during execution. This is achieved by having each individual portion of y be verifiable against x. Unfortunately, it has drawbacks we discuss later.

[3] Additionally, in proofs of useful (also known as catalytic) space (see [Pie19, Fis18, FBGB18, BDG17], and references therein), the cost of retrieval is the same as $\mathrm{cost}(f(x))$, which, as we just explained, must be greater than the cost of storing $(1 - \varepsilon_{\mathsf{space}}) \cdot |y|/(1 - \varepsilon_{\mathsf{hardness}})$ for time T_V; thus, a smaller $\varepsilon_{\mathsf{hardness}}$ translates into lower retrieval costs.

the cost of $f(x)$. For example, the original construction of Dziembowski et al. [DFKP15] has space ratio of at most $\frac{1}{512}$, the construction of Abusalah et al. [AAC+17] has space ratio that vanishes as $|y|$ grows, and the construction of Ren and Devadas has a space ratio of at most $\frac{1}{2}$, but only assuming that the temporary storage of P during the process of answering V is quite constrained.

There are only two known constructions that get at least one of these two parameters (i.e., $\varepsilon_{\mathsf{hardness}}$ and $\varepsilon_{\mathsf{space}}$) close to 0:

- The SDR construction of Fisch [Fis19] achieves arbitrarily small $\varepsilon_{\mathsf{space}}$; however, in that construction $\varepsilon_{\mathsf{hardness}}$ goes to 1 instead of to 0 as $\varepsilon_{\mathsf{space}}$ decreases.
- The $\mathsf{PoS}_{\mathsf{o}}$ construction of Pietrzak using a depth-robust graph [Pie19, Lemma 8] achieves arbitrarily small $\varepsilon_{\mathsf{space}}$ and $\varepsilon_{\mathsf{hardness}}$, but at an impractically high cost. Specifically, the computational complexity of f, per bit of y, grows as $(\varepsilon_{\mathsf{space}} \cdot \varepsilon_{\mathsf{hardness}})^{-3}$. And the computational complexity of the initialization protocol between P and V (in which P commits to y and proves that it was computed almost correctly) grows as $(\varepsilon_{\mathsf{space}} \cdot \varepsilon_{\mathsf{hardness}})^{-4}$. The parameters are believed to be impractical even for $\varepsilon_{\mathsf{space}} = \frac{1}{2}$ and $\varepsilon_{\mathsf{hardness}} = \frac{3}{4}$ (see [Fis19, Section 1.1]).

This, except for one construction of impractical (octic!) complexity, the problem originally posed by Dziembowski et al. [DFKP15] remains open: how to build a proof of space that requires the adversary to use almost as much storage or do almost as much work as the honest parties?

1.2 Our Contribution

We present a new proof-of-space construction that achieves arbitrarily small $\varepsilon_{\mathsf{space}}$ and $\varepsilon_{\mathsf{hardness}}$ while dramatically improving on the complexity of the depth-robust-based construction in [Pie19]. In our protocol, the computational complexity of f (per byte of y) grows as

$$\tilde{O}\left(\frac{1}{\varepsilon_{\mathsf{hardness}}}\right),$$

and the complexity of the initialization protocol between P to V grows as

$$\tilde{O}\left(\frac{1}{\varepsilon_{\mathsf{hardness}}}\left(\frac{1}{\varepsilon_{\mathsf{hardness}}}+\frac{1}{\varepsilon_{\mathsf{space}}}\right)\right).$$

We thus reduce the dependence on $\varepsilon_{\mathsf{space}}$ and $\varepsilon_{\mathsf{hardness}}$ from octic to nearly quadratic. We also pay attention to constants, showing the practicality of our construction.

In Section C, we show how our results improve concrete parameters for the currently deployed proof of space on the Filecoin blockchain. Without any change to the construction, our new analysis improves r_{hardness} by a factor of 11, from 1/55 to 1/5.

1.3 Technical Overview

We now elaborate on the construction and analysis.

The Model of Computation All known constructions of proofs of space are in the random oracle model; let H denote this random oracle. Queries to H are assumed to be atomic; space is measured in the number of H outputs stored, and time is measured in the number of queries to H. Like most constructions of proofs of space our construction, which we call SPR for reasons to be explained shortly, uses f that is computed via a directed acyclic graph G with a single source, as follows. Each node in G is labeled with the output of H applied to the labels of the node's predecessors (and the node's index in the graph, for uniqueness); the source is labeled with x and the labels of nodes near the sink(s) become y. (The only exception to this construction approach is the construction of [AAC$^+$17], which is based on computing tables of random functions and sorting them by output value.)

Like the original proof-of-space work by Dziembowski et al. [DFKP15], we assume that a malicious prover P^* stores whole labels of some of the nodes in G. Labels are treated atomically in our model; that is, we assume that the adversary does not attempt to save space by storing partial labels or functions of labels such as linear combinations, etc.[4] Then, when, during execution, a query from V asks P^* for the label of some node v in y, P^* needs to compute this label from the labels stored. This problem corresponds to the following pebbling game on G: given pebbles on some nodes on G (the ones with stored labels), work to place a pebble onto v; you are allowed to place a pebble onto a node when all of its predecessors already have pebbles. The complexity of a computation corresponds to the total number of pebbles placed in order to reach v.

A malicious prover can cheat somewhat when initially computing $y = f(x)$, by labeling some nodes with incorrect, easy to remember (e.g., pseudorandom), values. The initialization protocol between P and V ensures that this cheating cannot cover too many nodes: P commits to the labeling of G; V asks P to reveal the labels of some random nodes together with their predecessors and verifies that the computation of those labels was locally correct (the proof can be made noninteractive using the Fiat-Shamir heuristic). After $O(\frac{1}{\delta})$ queries, V can be sure that P did not cheat on more than δ fraction of the nodes; thus, the complexity of the initialization protocol is proportional to the graph degree times $\frac{1}{\delta}$. For P, incorrectly computed nodes correspond to additional pebbles on G, because they are already known without any work. Pebbles of this "cheating" type are called "red" [DFKP15], in contrast to pebbles that correspond to storage, which are "black".

When a node v is queried by V during execution, the prover has to compute the labels of all nodes that have unpebbled paths to v; thus, our goal is to a prove a high lower bound on the number of such nodes, no matter how the red and black pebbles are placed. Such nodes make up the *footprint* of v.

Most Relevant Prior Constructions The construction of [Pie19] mentioned above uses an (e, d)-depth-robust graph G for the computation of f. Such a graph guarantees

[4] It is of course reasonable to consider a malicious prover who stores information other than whole random oracle outputs — for example, functions of those outputs. For the case of parallel time discussed at the end of this section, it is known that this adversary is no more powerful than an adversary who treats random oracle outputs atomically as in our model — see [Pie19, Sections 5, 7]. Unfortunately, no proof of this fact is known for sequential time; in particular, finding such a proof for our construction remains an open problem. We are thus limited to this weaker model of the adversary.

the existence of an unpebbled path covering d fraction of the nodes even when e fraction of the nodes are pebbled. y consists of the labels of the topologically last portion of the graph. A relatively simple analysis shows that to get $\varepsilon_{\mathsf{space}}$ and $\varepsilon_{\mathsf{hardness}}$ approach 0, $e + d$ have to approach 1 (so-called "extreme" depth robustness), which results in the graph degree having to grow as $(\varepsilon_{\mathsf{space}} \cdot \varepsilon_{\mathsf{hardness}})^{-3}$ (per the analysis in [ABP18],[EGS75, Lemma 1]). To make the analysis go through, δ has to shrink as $O(\varepsilon_{\mathsf{space}} \cdot \varepsilon_{\mathsf{hardness}})$, and thus the verification protocol cost grows as $(\varepsilon_{\mathsf{space}} \cdot \varepsilon_{\mathsf{hardness}})^{-4}$.

The SDR ("Stacked Depth-Robust") construction of [Fis19] avoids the impracticality of extreme depth-robustness by using constant depth-robustness combined with expansion. It starts with a $(0.2, d)$ depth-robust graph for an arbitrary d. Take such a graph on n nodes and sort it topologically from left to right. Stack ℓ copies of this graph and connect consecutive layers using a (particular) bipartite expander directed down (the idea of using stacked expanders for proofs of space first appeared in [RD16]). The labels of the bottom layer constitute y. The proof [Fis19] shows that if the adversary saves $\varepsilon_{\mathsf{space}}$ fraction of the storage, then there are $\frac{\varepsilon_{\mathsf{space}}}{2} \cdot n$ nodes at the bottom layer such that each, in its footprint, contains an unpebbled path of length dn in some layer, as long as there are enough layers — specifically, as long as ℓ grows as $\log \frac{1}{\varepsilon_{\mathsf{space}}}$. Thus, $r_{\mathsf{hardness}} = d/\ell$, and $\varepsilon_{\mathsf{hardness}} = 1 - r_{\mathsf{hardness}}$ goes to 1 instead of 0.

Our Construction Our construction SPR ("Stacked Predecessor-Robust") is a slight generalization of SDR. The technical heart of this paper is in our new analysis. Recall that our goal is to show that $\varepsilon_{\mathsf{hardness}}$ can go to 0. We do so by showing that for a $\varepsilon_{\mathsf{space}}/2$ fraction of the bottom-layer nodes the following is true: they each have a footprint that covers almost the entire graph.

Specifically, we show a much better estimate of footprint sizes than the proof of [Fis19] implies. We do so in three main steps.

- First, we generalize the approach of [Fis19] to work for arbitrary depth-robust (actually, predecessor-robust) and expander graphs (rather than the specific ones used in [Fis19]), and develop new techniques to give tight bounds on exactly when the special layer with footprint dn occurs. The basic idea of the argument is simple: expander graphs make footprints grow as you go up, but pebbles reduce the growth. Eventually the pebbles run out, because the malicious prover P^* has only $(1 - \varepsilon_{\mathsf{space}})n$ black ones and some red ones. Getting a proof that can account for all possible allocations of pebbles chosen by P^* is where the technical difficulties lie.

- Second, unlike prior work, we consider how this footprint grows above the special layer. Again, getting a proof that can account for all possible allocations of pebbles chosen by P^* is where the technical difficulties lie. We show that if P^* does not have enough pebbles, the footprint will contain almost all the nodes on almost all the layers above, which is good enough for our purposes. But this argument alone is insufficient, because P^* may actually have enough pebbles to prevent the footprint above the special layer from growing. This brings us to the next step.

- The third part of our proof is to show that there are many special layers containing a footprint of size at least dn, and that pebbles can be used to prevent only a few

these footprints from growing in layers above; eventually, one of them will grow to fill up almost all the layers above it.

Technically, our construction can be viewed as amplifying predecessor robustness of graphs. Predecessor robust graphs are directed acyclic graphs in which every set of nodes of relative size π contains a large single-sink induced subgraph of relative order α_π (i.e., using α_π fraction of the nodes of the original graph and all the edges between them). We start with an n-node predecessor-robust graph for some constant π and α_π. We layer $\ell = O(\frac{1}{\varepsilon_{\mathsf{hardness}}} \log \frac{1}{\varepsilon_{\mathsf{space}}})$ copies of it using constant-degree expanders, growing the size by a factor of ℓ, but increasing the degree only additively by a constant (specifically, by 8). The properties of the graph we build can be phrased in terms of extreme predecessor robustness: any set of nodes of relative size $\frac{(\ell-1+\varepsilon_{\mathsf{space}})n}{\ell n} = 1 - O(\frac{(1-\varepsilon_{\mathsf{space}})\varepsilon_{\mathsf{hardness}}}{\log 1/\varepsilon_{\mathsf{space}}})$ (corresponding to the set of unpebbled nodes) has a single-sink induced subgraph (corresponding to the footprint of a bottom-layer node) of relative order $1 - \varepsilon_{\mathsf{hardness}}$, which is arbitrarily close to 1. Moreover, for proofs of space we require (and achieve) more: there are $\frac{\varepsilon_{\mathsf{space}}}{2} \cdot n$ such induced subgraphs with distinct sinks specifically on the bottom layer (these are the nodes that, when queried by V, will cause P^* a lot of work).

The Cost of Initialization The parameter ℓ is the ratio of the number of node labels computed during the computation of $y = f(x)$ (there are $\ell \cdot n$ of them) to the number of node labels stored as y (there are n of them). The PoS_o construction of [Pie19, Lemma 8] (in fact, any construction[5]) requires $\ell \geq (1 - \varepsilon_{\mathsf{space}})/\varepsilon_{\mathsf{hardness}}$. Even though our construction SPR has somewhat higher ℓ and thus more nodes to compute, the initialization in SPR is much faster because the degree of each node is much smaller. Moreover, SPR initialization can be completed using only $2 \cdot |y|$ memory (because each layer depends only on the previous one, and proofs of correctness can also be computed using just two layers at a time), while PoS_o requires the entire graph to be stored in memory during initialization, because long-range dependencies are essential for depth robustness [AGK+18, Section 3]. See Sections 2.2 and 8 for a more detailed analysis.

On Parallel vs. Sequential Time The original proofs of space definition [DFKP15], as well as the works of [AAC+17] and [RD16], consider the total computational effort (also known as sequential time) when measuring computational hardness of cheating by the malicious prover. We do so as well. In practice, this measure corresponds, roughly, to the cost of computation when cheating (to discourage cheating, this cost ought to be lower than the cost of storage for the time interval T_V between executions). We emphasize, however, that the works of [Pie19] and [Fis19] also consider a stronger notion: parallel time, i.e., the amount of computational steps required even if each step can use unbounded parallelism. In practice, this measure corresponds, roughly, to latency; to prevent cheating, the verifier should time out if the response does not come fast enough. This measure has the advantages of not making any assumption on the

[5] If there are fewer than $n(1 - \varepsilon_{\mathsf{space}})/\varepsilon_{\mathsf{hardness}}$ nodes in G, then an adversary who stores $n(1 - \varepsilon_{\mathsf{space}})$ nodes can simply recompute the rest of G during the execution protocol, and the cost will be less that $|G| - n(1 - \varepsilon_{\mathsf{space}}) = |G| - |G|\varepsilon_{\mathsf{hardness}} = (1 - \varepsilon_{\mathsf{hardness}})|G|$.

rationality of the adversary and not requiring cost estimation (which is uncertain in practice), though it still requires latency estimation. Both measures have found use in practice: the prevention of cheating in Chia is based on monetary cost [CP19, Section 2.2.2], while the prevention of cheating in Fielcoin is based on a mix of cost and latency [GN23, Section 1.2].

Extending our results to parallel time (i.e., latency) remains an open question. Technically, whereas sequential time corresponds to the total footprint size of a queried node v, parallel time corresponds to the longest path within that footprint. The only construction to achieve small $\varepsilon_{\mathsf{space}}$ and $\varepsilon_{\mathsf{hardness}}$ for parallel time is the aforementioned construction of [Pie19], which incurs impractically high cost due to the need for extremely depth-robust graphs, for which the only known constructions require very high degree. Extending our results to parallel time would also automatically get us provable security against a broader class of adversaries; namely, adversaries who are not limited to treating random oracle outputs atomically (see Footnote 4 above).

2 Definition and Construction

2.1 The Graph SPR

Please refer to Section 1.3 for the explanation of the construction; here we only fill in the details.

Recall that our construction SPR is a generalization of the SDR construction by Fisch [Fis19]. In both, $y = f(x)$ is computed by labeling a single-source directed acyclic graph G of ℓ levels of n nodes each. The label of the source node is x, the label of each node except the source is computed by hashing the labels of its predecessors, and the labels of the entire bottom level are y. We follow the numbering in [Fis19]: the top level is 1, the bottom is ℓ, with edges going left-to-right in each level (for depth-robustness or predecessor-robustness per level) and down from level i to level $i+1$ (for expansion when going back from lower levels to upper levels). Note that this level numbering can be confusing, as most of our arguments go bottom-to-top by induction, and thus induction goes down in natural numbers as it goes up levels.

SDR requires the following depth-robustness guarantee: any set of $0.8n$ nodes in a given level has a horizontal path of length dn. This guarantee needs to apply to all levels. In SPR, we relax the depth-robustness guarantee. We parameterize SPR by both ℓ and ℓ_{pr}. Of the ℓ total levels, only the lower ℓ_{pr} need to have the following guarantee, called *predecessor robustness* in [AdNV17]: on each of the lower ℓ_{pr} levels, any set of $\pi \cdot n$ nodes contains a single-sink induced subgraph with $\alpha_\pi \cdot n$ nodes (this guarantee is implied by depth-robustness, as a path is, in particular, a single-sink subgraph). In contrast to SDR, which is analyzed specifically for $\pi = 0.8$, our construction works for almost any constant π and α_π. There is a mild technical condition that relates π to the behavior of the expander; see Condition 2 in Section 4.1. The levels above the lowest ℓ_{pr} levels need no horizontal edges, except level 1, which needs an edge from the leftmost (source) node to every node.

In both SDR and SPR, the edges from level i to level $i+1$ form an expander when viewed backward; that is, a set of nodes on level $i+1$ of size αn has $\beta(\alpha) \cdot n$ predecessors

on level i, where $\beta(\alpha) > \alpha$, and the ratio $\beta(\alpha)/\alpha$ is greater than some constant for sufficiently small values of α. In contrast to SDR, whose analysis in [Fis19] is tightly tied to the specific function β (from the degree-8 Chung expander; see Appendix A), most of our analysis works for general β. For most of the analysis we require only the following:

Condition 1. $\beta(\alpha)$ is a continuous, monotonically increasing strictly concave function on $[0, 1]$, with $\beta(0) = 0$, $\beta(1) = 1$, and $\beta(\alpha) > \alpha$ for all $\alpha \in (0, 1)$.

We instantiate our vertical graphs with degree-8 Chung expanders only at the end of the proof, to measure ℓ. Other expanders would also work.

We emphasize that SDR is a special case of SPR, and our analysis works for SDR as well.

2.2 Initialization

As in all proof of space schemes (except the aforementioned [AAC+17]), initialization starts by having P compute the labeling of G, commit to the labels using a Merkle tree or another vector commitment, and send the commitment to V. To ensure the computation is approximately correct, V queries some number of randomly chosen nodes, whose labels P reveals together with the labels of their predecessors; V verifies that the decommitments are correct and that the label of each requested node is correctly computed from its predecessors. For our construction, initialization will assure V with probability $1 - e^{-\lambda}$ that the fraction of incorrectly computed labels on each layer is at most δ. This will require λ/δ queries per layer. The cost of initialization is thus $O\left(\frac{\ell_{\mathsf{pr}} \cdot d_{\mathsf{pr}} + \ell \cdot d_{\mathsf{exp}}}{\delta}\right)$, where d_{pr} and d_{exp} are the degrees of the predecessor-robust graph and expander, respectively.

Note that during initialization, P needs to have $2 \cdot |y|$ memory (i.e., double the long-term storage), because P needs to store just two layers: the current one being computed and the previous one.[6] The proofs of correctness (per-layer Merkle trees and openings of random nodes and their predecessors) can also be computed using just two layers at a time. Alternatively, if P has $\ell \cdot |y|$ memory, then the proof of correctness can be sped up somewhat by using a single Merkle tree on columns of nodes: because the graph is almost the same layer-to-layer, entire columns of nodes can be queried and decommitted at once.[7]

From now on, we will assume that initialization has succeeded: that is, we assume V has accepted initialization, the probability $e^{-\lambda}$ event that P^* was not caught cheating has not happened, and thus at most a δ fraction of each layer is incorrect. Nodes with incorrect labels will be said to have red pebbles on them.

[6] Ren, Devadas, and Fisch [RD16, Fis19] suggest a localization procedure to eliminate inter-layers edges that go right, in order to reduce the memory needed during initialization to just $|y|$; but we do not see how to make our proof of security go through if this procedure is applied

[7] It is also possible to use $\ell\lambda/\delta$ challenges for the entire graph G to guarantee that at most δ/ℓ fraction of the entire graph is incorrect, which would in particular imply the per layer guarantee of δ, but the per-layer approach is more efficient, because working with entire columns means that there are only λ/δ decommitments.

After initialization the honest P stores the n labels of the bottom layer of the graph. A malicious P^* stores the labels of any $(1-\varepsilon_{\mathsf{space}})\cdot n$ nodes; these nodes will be said to have black pebbles on them.

2.3 Execution and Security

During the execution, V queries a bottom-layer node, and P decommits its label. A malicious P^* must place new pebbles in order to find the label of P; recall that a pebble can be placed on a node only if all of its predecessors have pebbles.

Our definition of security is in the graph pebbling model, following [DFKP15, RD16]. Note that different definitions of proofs of space highlight different parameters in parentheses; we avoid the positional parenthetical notation to avoid confusion.

Definition 1. Let N be the number of nodes in G and n be the number of nodes in the output y. We will say that a proof of space in the pebbling model has space gap $\varepsilon_{\mathsf{space}}$, hardness gap $\varepsilon_{\mathsf{hardness}}$, and single-query catching probability p_{hard} if the following holds: assuming initialization succeeded, with probability at least p_{hard} over the random choice of a queried node, a cheating prover P^* who stores at most $(1-\varepsilon_{\mathsf{space}})\cdot n$ black pebbles before the query is issued must place pebbles onto $(1-\varepsilon_{\mathsf{hardness}})\cdot N$ nodes[8] in order to place a pebble onto the queried node.

Note that V can query $\lambda/p_{\mathsf{hard}}$ nodes to increase the probability from p_{hard} to $1-(1-p_{\mathsf{hard}})^{\lambda/p_{\mathsf{hard}}} > 1-e^{-\lambda}$ (however, the work of P^* might not grow above $(1-\varepsilon_{\mathsf{hardness}})\cdot N$, as it might be shared among all the queried nodes). Note also that p_{hard} cannot exceed $\varepsilon_{\mathsf{space}}$ (because the $1-\varepsilon_{\mathsf{space}}$ fraction of y may simply be stored); we achieve near-optimal $p_{\mathsf{hard}} = \varepsilon_{\mathsf{space}}/2$.

3 Main Result and Proof Overview

Theorem 1. *For any $\varepsilon_{\mathsf{space}} > 0$ and $\varepsilon_{\mathsf{hardness}} > 0$, there is a setting of parameters ℓ, ℓ_{pr}, δ, and n in the SPR construction that achieves space gap $\varepsilon_{\mathsf{space}}$, single-query catching probability $p_{\mathsf{hard}} \geq \varepsilon_{\mathsf{space}}/2$, and hardness gap $\varepsilon_{\mathsf{hardness}}$, such that*

- *The cost of computing $f(x)$, per bit of y, is*

$$O\left(\frac{1}{\varepsilon_{\mathsf{hardness}}}\cdot\log\frac{1}{\varepsilon_{\mathsf{space}}}\right).$$

- *The cost of the initialization protocol is*

$$\tilde{O}\left(\frac{1}{\varepsilon_{\mathsf{hardness}}}\cdot\left(\frac{1}{\varepsilon_{\mathsf{hardness}}}+\frac{1}{\varepsilon_{\mathsf{space}}}\right)\right).$$

[8] We define $\varepsilon_{\mathsf{hardness}}$ in terms of nodes pebbled rather than edges traversed. If the degrees of nodes are similar, it does not make much of a difference. We could, instead, redefine it in terms of edges traversed, which would account for the fact that costs of hashing are roughly proportional to the input length; this would make accounting messier, but would not change our main result of achieving $\varepsilon_{\mathsf{space}}$ and $\varepsilon_{\mathsf{hardness}}$ arbitrarily close to 0. It is also possible to use duplicate label inputs to H simply to make computation time at each node the same without increasing the degree (as is done in the Filecoin implementation [Pro23]).

.

The rest of the paper is dedicated to proving this theorem. We start by providing a proof overview.

Given a set S of nodes, let weight $wt(S)$ denote $|S|/n$.

Definition 2. A path is *unpebbled* if none of its nodes (including beginning and end) have pebbles. For a node v, its *footprint* is the set of nodes that have unpebbled paths to v. If v itself is pebbled, its footprint is empty. For a set of nodes, its footprint is the union of the footprints of its elements.

Fix a set weight ζ, with $1 - \varepsilon_{\mathsf{space}} + \delta < \zeta < 1$.

Definition 3. Call a level b *fertile* if for every subset S of the bottom level with $wt(S) \geq \zeta$, the footprint of S on level b has weight at least π. That means that the predecessor robustness guarantee applies to the footprint of S on level b, so the footprint on level b has a single-sink induced subgraph of weight α_π.

3.1 Summary of the SDR Proof from Fisch [Fis19]

Our goal is to show that sufficiently many nodes in the bottom level have sufficiently large footprints. We don't know how to do that using only expansion arguments (i.e., vertical edges), because we can't prove that an average single node in the bottom level expands much as we go up.

The proof in [Fis19] first uses the expansion argument on a *set* of nodes to prove that it expands, and then uses *horizontal* edges to prove that even a single node at the bottom will depend on many nodes in a given level. Specifically, the proof proceeds as follows (substituting predecessor-robustness for depth-robustness):

1. **Expansion to get a large footprint of a large set.** Prove, using vertical edges and expansion arguments, that there exists a fertile level. At its core, the argument is relatively simple: any bottom-level set S of weight ζ expands to the next level via β, pebbles reduce this expansion, and you repeat. Eventually pebbles run out and you win.

   The argument is suboptimal because $\beta(\alpha)$ for the specific degree-8 Chung expander used in [Fis19] is a messy function, and the proof uses its piecewise-linear approximation to reach $\pi = 0.8$. In Section 5, we replace this argument with one that works for a general π and a general β using its global properties from Condition 1; this improved argument gives better results (i.e., the fertile level is lower) even for the specific expander in [Fis19] (see Appendix C.1).

2. **Predecessor robustness to get a single-sink footprint.** By the predecessor robustness property, the level-b footprint of any bottom-level set S contains a single-sink induced subgraph T of relative order α_π.

3. **Single-sink graphs to go from collective to individual footprints.** At least *one* node in S depends on the sink of T, and therefore the individual footprint of that one node contains all of T and is thus of size α_π (note that in SDR, as

opposed to SPR, T is a chain because of depth robustness, which implies that pebbling this one node takes time α_π even with unbounded parallelism).

4. **Simple counting to get many nodes with large footprints.** Because the above holds for *every* S of weight ζ on level ℓ, there are at least $(1-\zeta)\cdot n$ nodes at level ℓ whose footprint at level b contains a graph T of weight α_π (else, all the nodes that don't satisfy this condition form a set S that contradicts the previous three steps).

3.2 Main Idea of the Improvement

Our proof that footprint size is $(1-\varepsilon_{\mathsf{hardness}})N$ for any $\varepsilon_{\mathsf{hardness}} > 0$ proceeds in the same steps as outlined above, but with the addition of a new step after Step 3 above:

3.5 **An individual footprint on a fertile level expands in levels above.** The single-sink graph T has a footprint T' that is of weight $(1-\varepsilon_{\mathsf{hardness}}/2)$ on almost every level above b.

Applying Step 4 to T' instead of T, and bounding the fraction of levels with insufficient footprints, we get that there are at least $(1-\zeta)\cdot n$ nodes at the bottom level whose footprint is of size $(1-\varepsilon_{\mathsf{hardness}})N$

To make step 3.5 work, we will need to argue that above b, there are not enough pebbles to kill this expansion of T. Unfortunately, that is not the necessarily the case, because α_π may be quite small and there may be a lot of pebbles left.

At its core, the argument will be as follows. Each infertile level costs the adversary some black pebbles, because S wants to expand, and it costs pebbles to keep this expansion in check (Section 5). This bounds the number of infertile levels. Each fertile level has an unpebbled set T that wants to grow. We characterize the minimum footprint of such a set in Section 6, where we use concavity of β to prove that to minimize the footprint weight, all the pebbles should be placed on the level directly above T.

The challenge is that the adversary has enough pebbles to completely prevent the growth of a single fertile level. Moreover, some black pebbles can be used to stop several fertile levels at once. In Section 7, we show that, despite this ability, for *each* fertile level that is prevented from growing, the adversary has to use some quantity of black pebbles. We then show that eventually some fertile level's footprint will outgrow the number of black pebbles that the adversary can use above it (the main insight here is to look at the gap between the footprint and available pebbles, thus reducing two variables to one). Once a fertile level's footprint outgrows the number of available pebbles, we can lowerbound the rest of the footprint, no matter how the pebbles above are distributed.

This argument shows that if we carefully choose a fertile level in Step 1, we will be done. We fill in the quantitative details in Section 8.

4 Proof Notation and Basic Notions

We recap notation used above and introduce some new notation.

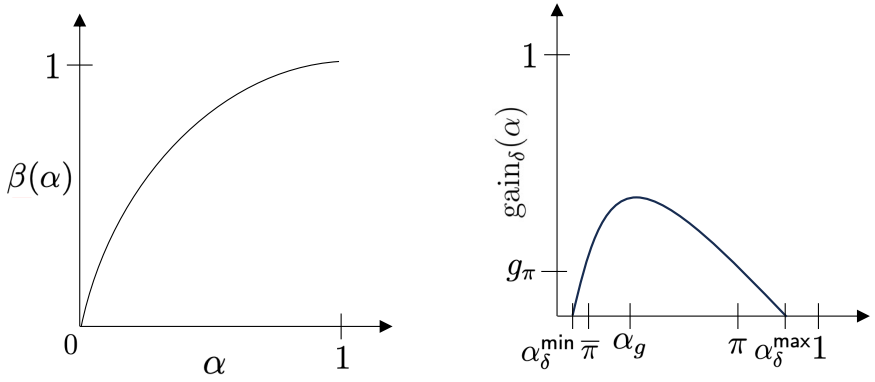

Figure 1: Generic β and $gain_\delta$ (see Figure 2 in Appendix A for the Chung expander).

4.1 Graph, Weights, Gains

The number of nodes at each level is n, and the total number of levels is ℓ, of which the lower ℓ_{pr} have horizontal edges to ensure predecessor robustness — i.e., to ensure that any subset of weight at least π has single-sink subgraph of weight α_π. Given a set S of nodes, let weight $wt(S)$ denote $|S|/n$. The levels are connected via an expander so that a subset of weight α on level i has $\beta(\alpha)$ predecessors on level $i-1$, with β satisfying Condition 1; let $gain(\alpha) = \beta(\alpha) - \alpha$, $\beta_\delta(\alpha) = \beta(\alpha) - \delta$, and $gain_\delta(\alpha) = gain(\alpha) - \delta$, where δ is maximum per-level weight of red pebbles.[9]

We prove the following standard set of facts in Appendix B.

Fact 1. *The function gain is strictly concave on the interval $[0,1]$, with $gain(0) = gain(1) = 0$. There is a value $0 < \alpha_g < 1$ that maximizes gain. The function gain (and therefore also $gain_\delta$) is monotonically increasing on inputs from 0 to α_g and monotonically decreasing on inputs from α_g to 1.*

We assume $\delta < gain(\alpha_g)$ (because we get to choose δ) and let $[\alpha_\delta^{\min}, \alpha_\delta^{\max}]$ denote the interval in which $gain(\alpha) \geq \delta$ (i.e., $gain_\delta(\alpha) \geq 0$). We do not care about expansion guarantees outside of this interval (thus, we can set n large enough so that expansion guarantees hold on the interval once the constants $\alpha_\delta^{\min}, \alpha_\delta^{\max}$ are fixed). Note that $\alpha_\delta^{\min} < \alpha_g < \alpha_\delta^{\max}$.

Condition 2. We assume $\pi > \alpha_g$.

If this condition does not hold, decrease δ and/or increase π until it does, which won't hurt predecessor robustness (but may change the constants in Theorem 1). Let

[9] One of the technical challenges in the proof is having to deal with δ red pebbles at every level, which means the total number of red pebbles can easily exceed n. If we had a small upper bound on the total number of red pebbles, we could just add them to the black pebbles, as long as the total was less than n. This would simplify the proofs and improve the quantitative bounds, but require more effort during initialization, as explained in Footnote 7.

$g_\pi = gain_\delta(\pi)$, and let $[\overline{\pi}, \pi]$ denote the interval in which $gain(\alpha) \geq g_\pi$. Note that $\overline{\pi} < \alpha_g < \pi$. Let $g_{\alpha_\pi} = gain_\delta(\alpha_\pi)$.

The following definition is used to measure how fast expansion grows over multiple layers when unimpeded by any black pebbles; our notation is on purpose analogous to logarithms.

Definition 4. Let $\beta\text{count}_x(y) = \min\{i : \beta_\delta^i(x) \geq y\}$, where β_δ^i denotes i repeated applications of β_δ.

Claim 1. *If $x > \alpha_\delta^{\mathsf{min}}$ and $y < \alpha_\delta^{\mathsf{max}}$, then $\beta\text{count}_x(y)$ is finite and is at most $\min(0, (y - x)/\min(gain_\delta(x), gain_\delta(y))$.*

Proof. If $x \geq y$, we are done, so assume $\alpha_\delta^{\mathsf{min}} < x < y < \alpha_\delta^{\mathsf{max}}$. Let $g = \min(gain_\delta(x), gain_\delta(y))$. We will be using Fact 1. Because $gain_\delta$ is strictly concave, $gain_\delta(x) > \min(gain_\delta(\alpha_\delta^{\mathsf{min}}), gain_\delta(\alpha_\delta^{\mathsf{max}})) \geq 0$ by Claim 15; same for $gain_\delta(y)$, so $g > 0$. Because $gain_\delta$ is concave, $gain_\delta(\alpha) \geq g$ for $x \leq \alpha \leq y$ by Claim 15, and therefore $\beta_\delta^i(x) \geq \min(y, x + i \cdot g)$, so $\beta\text{count}_x(y) \leq \min(0, (y - x)/g)$. □

Specifically for the degree-8 Chung expander (see Appendix A for details), expansion is rapid all the way from $x = \delta$ to $y = 1 - 3\delta$. We prove in Claim 10 that $\beta\text{count}_\delta(1 - 3\delta) \leq 2 \log 1/\delta$; and numerically, it is easy to calculate that $\beta\text{count}_\delta(1 - 2\delta) < 10$ for $\delta \geq 0.001$ (which seems to be a sufficiently small δ for all practical purposes).

4.2 Pebbles and Footprints

Let $\rho = 1 - \varepsilon_{\mathsf{space}}$ be the maximum total weight of black pebbles, ρ_i be the black pebble weight on level i, and $\rho_{i\ldots j}$ be black pebble weight on levels i through j, inclusive (regardless of whether $i \leq j$, i.e., $\rho_{i\ldots j} = \rho_{j\ldots i}$).

We will fix some parameter ζ to refer to the weight of a set S at level ℓ (it may help to think of $\zeta = (1 - \varepsilon_{\mathsf{space}})/2$, but in fact, we can pick any ζ as long as $\zeta > 1 - \varepsilon_{\mathsf{space}} + \delta + \alpha_\delta^{\mathsf{min}}$, $\zeta < 1$, and $\zeta < \delta + \alpha_\delta^{\mathsf{max}}$; a smaller ζ will increase the catching probability but decrease the footprint). Let $\zeta_\delta = \zeta - \delta$.

Recall Definition 2. If T is a subset of nodes at level j, define $f_j(T)$ to be the weight of the unpebbled part of T. For $i \leq j$, inductively define

$$f_i(T) = \max(0, \beta_\delta(f_{i+1}(T)) - \rho_i)\,. \tag{1}$$

Observe, by induction, that $f_i(T)$ is a lower bound on weight of the footprint of T at level i, because there are at least $\beta(f_{i+1}(T))$ parents of the footprint at level $i + 1$, and at most $\rho_i + \delta$ of those are pebbled (this is not even counting horizontal edges, if any). If $f_i = 0$, then for all $m \leq i$, $f_m = 0$, because $\beta_\delta(0) \leq 0$. Because f_i does not actually depend on the particular nodes in T, but only on f_j and j, we will also write $f_i(f_j, j)$, or even simply f_i when f_j and j are clear from the context.

Define the functions ϕ_T and ϕ_{f_j} as

$$\phi_T(\rho_j \ldots, \rho_1) = \phi_{f_j}(\rho_{j-1} \ldots, \rho_1) = f_j + \cdots + f_1$$

to provide a lower bound on the total footprint of T.

Our proof operates on the f and ϕ lowerbounds (rather than the actual footprints, which depend too much on the subsets chosen and the specific allocation of pebbles in each level). We will therefore start by studying the basic properties of these quantities.

We note that f and ϕ obey intuitive monotonicity constraints.

Claim 2. *For every $m \geq i$, f_i is monotone nonincreasing as a function of ρ_m and monotone nondecreasing as a function of f_j. The function ϕ is monotone: if $f'_j \geq f_j$ and $\rho'_i \leq \rho_i$ for each i, then $\phi_{f'_j}(\rho'_{j-1}, \ldots, \rho'_1) \geq \phi_{f_j}(\rho_{j-1}, \ldots, \rho_1)$. Moreover, adding a level at the end cannot decrease ϕ: $\phi_{f_j}(\rho_{j-1}, \ldots, \rho_1, \rho_0) \geq \phi_{f_j}(\rho_{j-1}, \ldots, \rho_1)$ for any ρ_0.*

Proof. The first sentence follows by monotonicity of β (see Condition 1). The second sentence can be proven by induction, as follows: using monotonicity of β, observe that the f_i values do not decrease when we change from f_j to f'_j or from ρ to ρ'. The third sentence follows from nonnegativity of f. □

4.3 Basic Facts about Bounding Footprints

The following simple claim will turn out surprisingly useful in lowerbounding footprints, because it will allow us to focus on the total amount of black pebbles rather than on their allocation to specific levels.

Claim 3. *Let T be a subset of level j and let $i \leq j$.*

$$f_i = \max\left(0, f_j + gain_\delta(f_j) + \cdots + gain_\delta(f_{i+1}) - \rho_{i\ldots j-1}\right) .$$

Proof. By induction on i starting at j and going down to 1. The base case is trivial. For the inductive case, note that if $f_i = 0$, then $f_{i-1} = 0$ because $\beta(0) = 0$ by Condition 1, and the formula in the claim also gives us 0 because $gain(0) = 0$ and $\rho_{i-1} \geq 0$. Else, $f_i = f_j + \sum_{k=i+1}^{j} gain_\delta(f_k) - \rho_{i\ldots j-1}$ by the inductive hypothesis, so

$$f_{i-1} = \beta_\delta(f_i) - \rho_{i-1} = f_i + gain_\delta(f_i) - \rho_{i-1} = f_j + \sum_{k=i}^{j} gain_\delta(f_k) - \rho_{i-1\ldots j-1} .$$

□

We generally will be interested in footprints that grow as we go up. The following claim allows us to rule out some situations in which the value of f decreases even without any black pebbles. This decrease can occur when f is too low (below $\alpha_\delta^{\mathsf{min}}$) or too high (above $\alpha_\delta^{\mathsf{max}}$), because then $gain_\delta$ is negative, and thus red pebbles alone are enough to decrease f. The claim shows, in part, that if f starts below $\alpha_\delta^{\mathsf{max}}$, it will always stay below $\alpha_\delta^{\mathsf{max}}$, essentially because $\beta_\delta(\alpha)$ cannot overcome the $\alpha_\delta^{\mathsf{max}}$ barrier.

Claim 4. *If for some m, $gain_\delta(f_m) > 0$, then for all $i \leq m$, either $f_i \leq \alpha_\delta^{\mathsf{min}}$ or $gain_\delta(f_i) > 0$. Moreover, if $\rho_{m-1\ldots i} = 0$, then $f_m < f_{m-1} < \cdots < f_i$.*

Proof. We will proceed by induction starting at $i = m$ and going down to $i = 1$. The base case is given. For the inductive step (going from f_i to f_{i-1}), consider the following cases that cover all the possibilities with $\alpha_\delta^{\mathsf{min}} < f_{i-1}$.

- if $\alpha_\delta^{\mathsf{min}} < f_{i-1} < f_i$, then $gain_\delta(f_{i-1}) > \min(gain_\delta(\alpha_\delta^{\mathsf{min}}), gain_\delta(f_i))$ by strict concavity of $gain_\delta$ (per Fact 1 and Claim 15), and $gain_\delta(f_i) > 0 = gain_\delta(\alpha_\delta^{\mathsf{min}})$ by the inductive hypothesis.
- if $\alpha_\delta^{\mathsf{min}} < f_i < f_{i-1}$, then, since $f_{i-1} > 0$, we know $f_{i-1} = \beta_\delta(f_i) - \rho_{i-1}$. Then $gain_\delta(f_{i-1}) = \beta_\delta(f_{i-1}) - f_{i-1} = \beta_\delta(f_{i-1}) - (\beta_\delta(f_i) - \rho_{i-1}) \geq \beta_\delta(f_{i-1}) - \beta_\delta(f_i) > 0$ by monotonicity of β_δ (Condition 1).
- If $\alpha_\delta^{\mathsf{min}} < f_i = f_{i-1}$ then $gain_\delta(f_{i-1}) = gain_\delta(f_i) > 0$ by the inductive hypothesis
- the case $f_i \leq \alpha_\delta^{\mathsf{min}} < f_{i-1}$ is impossible, because then $f_{i-1} \leq \beta_\delta(f_i) = f_i + gain_\delta(f_i) \leq f_i$, because $gain_\delta(f_i) \leq gain_\delta(\alpha_\delta^{\mathsf{min}}) \leq 0$ by Fact 1.

If, moreover, $\rho_{m-1\ldots i-1} = 0$, then $f_{i-1} \geq f_i + gain_\delta(f_i)$. Because $f_i \geq f_m$ by the inductive hypothesis and $f_m > \alpha_\delta^{\mathsf{min}}$ (by Fact 1, because we assume $gain_\delta(f_m) > 0$), we know $gain_\delta(f_i) > 0$ by the inductive hypothesis. Thus, $f_{i-1} > f_i$. □

5 Upperbounding the Number of Infertile Levels

In this section, we provide the theorem necessary for carrying out Step 1 in Section 3.1.

Recall the definition of fertile (Definition 3) and constraints on ζ (Section 4.2). The main idea for bounding the number of infertile levels is the following. Suppose some level b is infertile. That means there is a subset S of nodes on the bottom level ℓ of weight ζ whose level-b footprint is less that π. Note that $f_\ell = \zeta_\delta - \rho_\ell$ is a lower bound on the footprint of S on level ℓ, and $f_i(f_\ell, \ell)$ is a lower bound on the footprint of S on level i by Claim 2. It suffices to prove that the number of levels i with $f_i < \pi$ is small, because all infertile levels have this property. Consider two cases:

- If for all i, $f_i \geq \overline{\pi}$, then the gain $gain_\delta(f_i)$ of every infertile level i is at least g_π. Per claim Claim 3, every such gain has to be overcome with pebbles in order to keep the footprint from growing. Thus, every infertile level costs at least g_π in black pebbles, or else the footprint grows.
- If for some i, $f_i < \overline{\pi}$, most of the black pebbles must be at level i or below, per Claim 3, because $\overline{\pi}$ is small. The value of f will grow above i until it gets to π, and there are not many pebbles left to stop this growth, so as soon as the value of f reaches π, the remaining levels above will be fertile.

These cases essentially correspond to two possible adversarial strategies for placing pebbles: either keep as many f_i values as possible just below π and spend at least $g_\pi(\pi)$ black pebbles to keep f_{i-1} below π, or spend all the black pebbles on level ℓ to get a very small f_ℓ, which may cause a few lower levels to be infertile. In this section we show that the best adversarial strategy will not do much better than either of these two. Our bounds on the number of levels with $f_i < \pi$ are nearly tight, as we further discuss below.

As a result, we obtain the following theorem.

Theorem 2. *Assume $g_\pi > 0$ and $\alpha^{\mathsf{min}}_\delta + \rho < \zeta_\delta < \alpha^{\mathsf{max}}_\delta$. The number of infertile levels is less than*

$$\max\left(1 + \frac{\rho + \pi - \zeta_\delta}{g_\pi}, 1 + \beta\mathrm{count}_{\zeta_\delta - \rho}(\pi)\right),$$

and the first argument of max *is greater than the second whenever $\zeta_\delta - \rho \geq \overline{\pi}$.*

The rest of this section is dedicated to the proof of this theorem.

Proof. The following variant of Claim 3 specialized for the set S will be useful for us.

Claim 5. *Let $f_\ell = \zeta_\delta - \rho_\ell$. Assume $\alpha^{\mathsf{min}}_\delta + \rho < \zeta_\delta < \alpha^{\mathsf{max}}_\delta$. Then for all m (with $1 \leq m \leq \ell$)*

$$f_m(f_\ell, \ell) = \zeta_\delta + \mathit{gain}_\delta(f_\ell) + \cdots + \mathit{gain}_\delta(f_{m+1}) - \rho_{m\ldots\ell}\,,$$

and $\mathit{gain}_\delta(f_m) > 0$.

Proof. The intuition is that at every level, because gain_δ is positive below m, there are not enough black pebbles for f_m to go below $\alpha^{\mathsf{min}}_\delta$, and thus gain_δ will remain positive by Claim 4.

Formally, we proceed by induction on m from ℓ down to 1. For the base case, $0 < \alpha^{\mathsf{min}}_\delta < \zeta_\delta - \rho_\ell < \alpha^{\mathsf{max}}_\delta$, so $\mathit{gain}_\delta(f_\ell) > 0$. For the inductive case (going from m to $m-1$), observe that $f_m \geq \zeta_\delta - \rho_{m\ldots\ell}$ because $\mathit{gain}_\delta(f_i) > 0$ for $m \leq i \leq \ell$ by the inductive hypothesis. Therefore,

$$\begin{aligned}
f_{m-1} &\geq f_m + \mathit{gain}_\delta(f_m) - \rho_{m-1} && \text{by Claim 3}\\
&\geq f_m - \rho_{m-1} && \text{by the inductive hypothesis}\\
&\geq \zeta_\delta - \rho_{m\ldots\ell} - \rho_{m-1} && \text{as shown about } f_m \text{ above}\\
&\geq \zeta_\delta - \rho > \alpha^{\mathsf{min}}_\delta\,.
\end{aligned}$$

Thus, f_{m-1} is positive and the formula follows by Claim 3; since $f_{m-1} > \alpha^{\mathsf{min}}_\delta$, we have $\mathit{gain}_\delta(f_{m-1}) > 0$ by Claim 4. □

Theorem 2 now follows from Lemmas 1 and 2 below, once we remember that every infertile level must sastify $f_i < \pi$. Note that the first argument to the max is greater than the second when $\zeta_\delta - \rho \geq \overline{\pi}$ by Claim 1. □

5.1 The simpler case: when the footprints don't get too small

Lemma 1. *Assume $g_\pi > 0$ and $\alpha^{\mathsf{min}}_\delta + \rho < \zeta_\delta < \alpha^{\mathsf{max}}_\delta$. Let $m < \ell$ be some level. Assume that for all levels $i \geq m$, $f_i \geq \overline{\pi}$, and for at least $k > 0$ levels $i \geq m$, $f_i < \pi$. Then*

$$\rho_{\ell\ldots m} > \zeta_\delta - \pi + g_\pi \cdot (k-1)$$

and thus the total number of levels with $f_i < \pi$ is less than

$$1 + \frac{\rho - \zeta_\delta + \pi}{g_\pi}\,.$$

The bound in this lemma is tight if $\zeta_\delta \geq \pi$, because there is a matching adversarial strategy: spend $\rho_\ell > \zeta_\delta - \pi$ black pebbles on level ℓ and g_π black pebbles on every subsequent level until pebbles run out. If $\zeta_\delta < \pi$, then the adversary would have to spend more pebbles than stated in the bound, because the bound does not take into consideration higher gain $gain_\delta(\zeta_\delta) > g_\pi$ for the bottom level and a few levels above it. This makes a difference only if $\rho > \zeta_\delta$ is considerably smaller than π (i.e., the space gap is large).

Proof. There are k levels i at or below level m with $\overline{\pi} \leq f_i < \pi$, and for each of those, $gain_\delta(f_i) \geq g_\pi$. The gains of other levels are positive by Claim 5. Let m' be the highest level at or below level m with $\pi > f_{m'}$. Then by Claim 5

$$\pi > f_{m'} \geq \zeta_\delta + \sum_{i=m'+1}^{\ell} gain_\delta(f_i) - \rho_{\ell \ldots m'} \geq \zeta_\delta + (k-1) \cdot g_\pi - \rho_{\ell \ldots m} \,.$$

Rearranging the terms concludes the proof. □

5.2 The more complex case: small footprints

Lemma 2. *Assume $g_\pi > 0$ and $\alpha^{\mathsf{min}}_\delta + \rho < \zeta_\delta < \alpha^{\mathsf{max}}_\delta$. Assume for some level i, $f_i < \overline{\pi}$. Then the number of levels i with $f_i < \pi$ is at most $\beta\mathrm{count}_{\zeta_\delta - \rho}(\pi)$.*

This bound is tight up to one level, as the adversary has a matching strategy: place all ρ black pebbles on level ℓ; there will be at least $\beta\mathrm{count}_{\zeta_\delta - \rho}(\pi) - 1$ levels with $f_i < \pi$.

Proof. In order to use more intuitive language, in this proof we will slightly abuse notation and use the word "footprint" to denote f_i (even though it is merely a lower bound on the footprint of S) and the words "fertile" and "infertile" denote levels with $f_i \geq \pi$ and $f_i < \pi$, respectively.

Starting with some pebble allocation, we will proceed to rearrange the pebbles so at not to decrease the number of infertile levels. After all the rearranging is done, the black pebble weight will be all at the bottom level, except perhaps less than g_π on the highest infertile level. Since the lowest infertile level has footprint at most $\zeta_\delta - \rho$ and the second-to-highest infertile level k has footprint $f_k < \pi$, and there are no pebbles on levels $\ell - 1, \ldots, k$, the number of infertile levels is at most $\beta\mathrm{count}_{\zeta_\delta - \rho}(\pi)$.

The intuition is that packing more pebbles into a level with an already tiny footprint is best for the adversary, because the gain will be small, so the footprint will grow very slowly. Turning this intuition into a proof takes a sequence of carefully chosen steps.

No matter how the pebbles are arranged, every footprint is positive by Claim 5. Suppose level b is the lowest level with $f_b < \overline{\pi}$, and every level up to $m \leq b$ is infertile, while level $m - 1$ (if $m > 1$) is fertile.

First, if any level i is fertile and has pebbles above it, simply lower all the pebbles above it by one level. Let f'_i denote the new footprint at level i. Note that $f'_i = f_{i-1} - gain_\delta(f_i)$ is smaller than the old f_{i-1} (by Claim 5), and thus all the levels above i that were infertile will remain infertile, just one level lower (by monotonicity, Claim 2). Do so repeatedly until level ℓ is infertile and infertile levels continue, without gaps, until some level m.

Second, if the lowest level b with $f_b < \overline{\pi}$ is not ℓ, we know from Lemma 1 that $\rho_{\ell \ldots b} > \zeta_\delta - \pi + g_\pi \cdot (\ell - b) + \rho_b$. Move all the pebbles from levels $\ell - 1, \ldots, b$ down to level ℓ and let f'_ℓ denote the new footprint at level ℓ; $f'_\ell = f_b - gain_\delta(f_\ell) - \cdots - gain_\delta(f_{b+1}) < f_b - g_\pi \cdot (\ell - b)$, because $gain_\delta(f_i)$ for $i > b$ was at least g_π (because $f_i \in [\overline{\pi}, \pi]$). The new gain of each level up to b is less than g_π by induction (because $f_b < \overline{\pi}$), so the footprint at level b is at most f_b, and thus the footprints above level b have not increased by monotonicity, so the number of infertile levels has not decreased.

We can now assume that there are sufficient pebbles on level ℓ to cause $f_\ell < \overline{\pi}$ and that infertile levels continue without interruption until level m, with no higher infertile levels or black pebbles (if any are left, move them to m). If $m = \ell$, we are done, because $\beta\text{count}_{\zeta_\delta - \rho}(\pi)) \geq 1$, because $\zeta_\delta - \rho < \overline{\pi} < \pi$ by Claim 5. If $m = \ell - 1$, move all the pebbles form level $\ell - 1$ to level ℓ; this will decrease f_ℓ and therefore will decrease $gain(f_\ell)$ by Fact 1, because $f_\ell < \overline{\pi} < \alpha_g$, and therefore will decrease $f_{\ell-1}$, thus not decreasing the number of infertile levels. Thus, assume $m \leq \ell - 2$ for the rest of this proof.

We will describe an iteration of steps that reallocates pebbles. Each step will not decrease the number of infertile levels and will keep $f_\ell < \overline{\pi}$. We will always be able to take a step until $\rho_i = 0$ for all $m < i < \ell$ and $\rho_m < g_\pi$. At each step, we will either increase the number of levels i for which $\rho_i = 0$ without decreasing ρ_ℓ, or increase ρ_ℓ by at least g_π, and therefore the sequence of steps will be finite. At the end, we will have all pebbles on level ℓ, except at most g_π on level m.

At each step in the iteration, we do one of the following, specified in order of priority, unless none can be performed.

- Case 1. Suppose there exists a level $i \leq \ell - 2$ with $\rho_{i+1} + \rho_i \geq g_\pi$ and $gain_\delta(f_{i+1}) \geq g_\pi$. Move g_π of the pebbles from levels i and $i + 1$ to the bottom level ℓ and shift the pebbles from levels $\ell - 1, \ldots, i + 1$ up one level.

  Let f'_i denote the footprints after this step. Then $f'_\ell = f_\ell - g_\pi$, so $\beta_\delta(f'_\ell) = f_\ell - g_\pi + gain_\delta(f'_\ell) < f_\ell - g_\pi + gain_\delta(f_\ell) < f_\ell$ (where the first inequality follows by monotonicity of $gain_\delta$ below α_g, Fact 1; and the second by $f_\ell < \overline{\pi}$.). Note that $f'_{\ell-1} = \beta_\delta(f'_\ell) < f_\ell$, because there are no black pebbles left on level $b - 1$. Thus, by induction and monotonicity (Claim 2), for all $i \in [\ell - 1, m + 1]$, $f'_i < f_{i+1}$, so all levels up to $m+1$ remain infertile. Because level m now contains black pebbles that were formerly on level $m + 1$, as well as its own black pebbles, except for g_π ones that were moved, $f'_m = \beta_\delta(f'_{m+1}) - \rho_{m+1} - \rho_m + g_\pi < \beta_\delta(f_{m+2}) - \rho_{m+1} - \rho_m - g_\pi = f_{m+1} - \rho_m + g_\pi = \beta_\delta(f_{m+1}) - gain_\delta(f_{m+1}) - \rho_m + g_\pi = f_m - gain_\delta(f_{m+1}) + g_\pi \leq f_m$, so level m also remains infertile.

  For the rest of the cases, we assume no such level exists in this iteration.

- Case 2. Suppose there is a level $i < \ell$ level with $f_i < \alpha_g$. We want to show that $f_{i+1} < \alpha_g$. If $i = \ell - 1$, that is true because $f_\ell < \overline{\pi} < \alpha_g$. Else, suppose not. Since f_{i+1} is infertile and at least $\alpha_g > \overline{\pi}$, $gain(f_{i+1}) > g_\pi$, so $f_{i+1} + gain_\delta(f_{i+1}) \geq \alpha_g + g_\pi$, so $\rho_i > g_\pi$, but that contradicts the assumption in Case 1. Thus, $f_{i+1} < \alpha_g$.

  Move the ρ_i pebbles down from level i to level $i + 1$. This will reduce f_{i+1} and will change f_i from $\beta_\delta(f_{i+1}) - \rho_i = f_{i+1} + gain_\delta(f_{i+1}) - \rho_i$ to $\beta_\delta(f_{i+1} - \rho_i) = f_{i+1} + gain_\delta(f_{i+1} - \rho_i) - \rho_i$. By monotonicity of $gain_\delta$ (Fact 1) and the fact that

$f_{i+1} < \alpha_g$, this reduces f_i and therefore, by monotonicity of f_j (Claim 2), also reduces all f_j for $j < i$, thus not decreasing the number of infertile levels.

For the rest of the cases, we assume no such level exists in this iteration.

- Case 3. Let i be the highest level with $m < i < \ell$ for which there are any black pebbles, i.e., $\rho_i > 0$. If $i = \ell$ (or there is no such level at all), we are done. Note that $f_i \geq \alpha_g$, because otherwise we would have applied Case 2, and because $gain_\delta$ is positive by Claim 5, the same is true of $f_{i-1}, \ldots, f_{m+1}$.

  - Case 3a. Suppose $\rho_i \geq g_\pi$. Then $i = \ell - 1$ or $gain_\delta(f_{i+1}) < g_\pi$ (else Case 1 applies), so either way $gain_\delta(f_{i+1}) < g_\pi$, so $f_{i+1} < \overline{\pi} < \alpha_g$. We can move the pebbles down from level i to $i+1$, by the same argument as in Case 2.
  - Case 3b. Suppose $\rho_i < g_\pi$. For $m+1 \leq j \leq i$, $f_j \leq \pi - g_\pi$ (because $f_j = f_{j-1} - gain_\delta(f_j)$ and $f_{j-1} < \pi$ and $gain_\delta(f_j) > g_\pi$ because levels $j, j-1$ are infertile and $f_j \geq \alpha_g$). Move the ρ_i pebbles from level i to level $m+1$. This will increase f_j for $m+1 < j \leq i$, and therefore decrease their gains, so each f_j for $m+1 < j \leq i$ will increase by at most ρ_i, and thus will remain infertile because $\rho_i < g_\pi$. f_{m+1} will decrease because of the decrease in the gains below it and therefore f_j for $j \leq m$ will decrease by monotonicity (Claim 2). Thus, the number of infertile levels will not decrease.

    Now that these pebble are on level $m+1$, call their weight ρ_{m+1} instead of ρ_i and use f_{m+1} and f_m for the post-move footprints of the respective levels. If $\rho_{m+1} + \rho_m \geq g_\pi$, apply the same process as in Case 1 to move them to level b. We have thus created a new level with 0 black pebbles. Else, moving these pebbles from level $m+1$ to level m will not reduce the number of infertile levels, as we show in the next paragraph, and we will do so to create a new level with 0 black pebbles.

    Indeed, suppose otherwise. f_m and footprints of levels above m decrease, by the same argument as two paragraphs ago. Thus, if a new fertile level gets created by this move, then $\beta_\delta(\rho_{m+2}) \geq \pi$. But because m is infertile, we know

    $$f_{m+1} + gain_\delta(f_{m+1}) - \rho_m < \pi\,.$$

    Plugging in $\beta_\delta(f_{m+2}) - \rho_{m+1}$ for f_{m+1}, we get

    $$\beta_\delta(f_{m+2}) - \rho_{m+1} + gain_\delta(f_{m+1}) - \rho_m < \pi\,.$$

    Recalling that $\beta_\delta(\rho_{m+2}) \geq \pi$, we have

    $$\rho_{m+1} + \rho_m > gain_\delta(f_{m+1}) > g_\pi$$

    because level $m+1$ is infertile and $f_{m+1} \geq \alpha_g > \overline{\pi}$. This is a contradiction.

Thus concludes the proof of Lemma 2. □

6 Lowerbounding Footprints of Fertile Levels

In this section, we switch from thinking about per-level footprints of a set S of weight ζ at level ℓ to thinking about the *total* footprint of a set T at level b with that has unpebbled weight f_b, in order to carry out Step 3.5 from Section 3.2.

Recall that ϕ (defined in Section 4.2) is a lower bound on the total footprint. Naturally, the adversary's goal is to place black pebbles so as to minimize ϕ. While computing ϕ for specific input values is easy numerically, we wish to find a general lower bound on ϕ as a function of the total number of pebbles $\phi_{b-1\ldots1}$, without having to enumerate possible individual placements.

It may be intuitive to think that moving a pebble one level down always decreases ϕ, because growth stops earlier. It turns out that this intuition is not true in general (it is not difficult to build a counterexample for concrete parameters), because the value of f on the higher of the two levels may grow slightly as the pebble moves down, which will cause the values of f in the levels above it to also grow, compensating for the reduction. For example, for the parameters of Appendix C, $\phi_{0.2}(0, 0.7, 0) \approx 1.007$, while moving pebbles of weight 0.06 down increases it to $\phi_{0.2}(0.06, 0.64, 0) \approx 0.021$.

The main result of this section is the following theorem that shows that moving *all* black pebbles down results in the minimal possible ϕ.

Theorem 3. *Assume T is an unpebbled set at layer b of weight f_b. Assume $gain_\delta(f_b) > 0$ and let $\sigma = \beta_\delta(f_b) - \rho_{b-1\ldots1}$. Assume $\sigma > \alpha_\delta^{\mathsf{min}}$. Then*

$$\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) \geq \phi_{f_b}(\rho_{b-1\ldots1}, \underbrace{0, \ldots 0}_{b-2}) = f_b + \sum_{i=0}^{b-2} \beta_\delta^i(\sigma)\,.$$

Proof. The heart of the proof is the following Lemma 3. It says that moving all black pebbles one level down from the highest level with any black pebbles will decrease (or at least not increase) ϕ, as long as $gain_\delta(f_b) \geq 0$. This lemma, applied repeatedly $b-2$ times for $m = 1, 2, \ldots, b-2$, suffices for proving that the smallest ϕ is with all the black pebbles as low as possible. This implies the inequality. The equality follows simply by computing the value of f at each level; we need only to make sure we don't apply β_δ to negative numbers, which follows from $\beta_\delta(f_b) - \rho_{b-1\ldots1} \geq \alpha_\delta^{\mathsf{min}}$. □

Lemma 3. *Assume T is an unpebbled set at layer b of weight f_b and $gain_\delta(f_b) > 0$. Let $m = \min_i \rho_i > 0$ be the highest level with any black pebbles. If $m < b - 1$, then moving all these pebbles down one level will not increase ϕ. That is, for all b, f_b, and $\rho_m, \ldots, \rho_{b-1}$, the following holds as long as $gain_\delta(f_b) > 0$.*

$$\begin{aligned}&\phi_{f_b}(\rho_{b-1}, \ldots, \rho_{m+2}, \rho_{m+1}, \rho_m, \underbrace{0, \ldots 0}_{m-1})\\ \geq&\phi_{f_b}(\rho_{b-1}, \ldots, \rho_{m+2}, \rho_{m+1} + \rho_m, 0, \underbrace{0, \ldots 0}_{m-1})\end{aligned}$$

Before proving this lemma, we will prove the following simple claim.

Claim 6. *Suppose $f_i \le \alpha_g$. Then moving any black weight from level $i-1$ to level i will not increase ϕ.*

Proof. The footprint below level i will not change. Suppose the total weight of moved pebbles is $x \ge 0$. Then by Claim 3, f_i will decrease by x (but will not go below 0). Again by Claim 3, f_{i-1} will decrease by $gain_\delta(f_i) - gain_\delta(\max(0, f_i - x))$ (but not below 0), which is nonnegative because $gain_\delta$ is monotonically increasing below α_g (Fact 1). By Claim 2, none of the $f_{i-1}, \ldots, f_1$ will increase, and thus ϕ will not increase. □

Proof of Lemma 3. We will consider three different pebble arrangements:

- $\rho_{b-1}, \ldots, \rho_{m+2}, \rho_{m+1}, 0, \underbrace{0, \ldots 0}_{m-1}$ (with ρ_m completely removed)
- $\rho_{b-1}, \ldots, \rho_{m+2}, \rho_{m+1}, \rho_m, \underbrace{0, \ldots 0}_{m-1}$ (as in the left-hand side, with black pebbles of weight ρ_m on level m)
- $\rho_{b-1}, \ldots, \rho_{m+2}, \rho_{m+1} + \rho_m, 0, \underbrace{0, \ldots 0}_{m-1}$ (as in the right-hand side, with black pebbles of weight ρ_m moved to level $m+1$)

Denote the per-level footprint bounds (as computed via Equation (1) in Section 4.2) in the three cases by f_i, g_i, and h_i, respectively, and the totals f, g, and h. We need to prove that $g \ge h$. Because Claim 6 covers the case of $f_{m+1} = g_{m+1} \le \alpha_g$, it suffices to consider the case when $f_{m+1} = g_{m+1} > \alpha_g$.

The challenge in proving the desired result is that it may not necessarily be the case that $g_i \ge h_i$, because the g sequence has less time to grow to make up for the ρ_m pebbles, because ρ_m pebbles appear later in the sequence. The trick to this proof is to study how g_i recovers from ρ_m pebbles as compared to h_{i+1}.

The intuition is roughly this: placing pebbles on level $m+1$ causes a higher reduction in the footprint that placing the same pebbles on level m, because the function β is more sensitive on smaller inputs, and $f_{m+1} < f_m < f_{m-1} < \cdots < f_1$, so placing pebbles lower affects smaller inputs to β. Note that this intuition (and the result) no longer holds if there are pebbles at levels above m, because the f values are not necessarily increasing as we go up. We will now formalize this intuition.

Case 1: No 0s among footprints. It will be easier to first handle the case when all the f_i, g_i, and h_i values are nonzero, as this simplifies formula (1) to $f_i = \beta_\delta(f_{i+1}) - \rho_i$ (and similarly for g_i and h_i).

We need to prove that $g \ge h$. We will do so by proving that $f - g < f - h$: that is, placing pebbles on level m reduces ϕ less than placing pebbles on level $m+1$ does.

To compute $f-g$, observe that $f_i = g_i$ for $i > m$. Then $f_m - g_m = \rho_m$, $f_{m-1} - g_{m-1} = \beta_\delta(f_m) - \beta_\delta(f_m - \rho_m)$, and in general for $1 \le i < m$, $f_{m-i} - g_{m-i} = \beta_\delta^i(f_m) - \beta_\delta^i(f_m - \rho_m)$, where β_δ^i denotes β_δ applied i times. Thus,

$$f - g = \rho_m + \sum_{i=1}^{m-1} \beta_\delta^i(f_m) - \beta_\delta^i(f_m - \rho_m)\,.$$

To compute $f - h$, observe that $f_i = h_i$ for $i > m+1$. Then $f_{m+1} - h_{m+1} = \rho_m$, $f_m - h_m = \beta_\delta(f_{m+1}) - \beta_\delta(f_{m+1} - \rho_m)$, and in general for $1 \le i < m+1$, $f_{m-i+1} - h_{m-i+1} = \beta^i_\delta(f_{m+1}) - \beta^i_\delta(f_{m+1} - \rho_m)$. Thus,

$$f - h = \rho_m + \sum_{i=1}^{m} \beta^i_\delta(f_{m+1}) - \beta^i_\delta(f_{m+1} - \rho_m) > \rho_m + \sum_{i=1}^{m-1} \beta^i_\delta(f_{m+1}) - \beta^i_\delta(f_{m+1} - \rho_m)$$

where the inequality follows from the fact that β_δ is monotonically increasing (Condition 1), so β^i_δ is monotonically increasing, and $\rho_m > 0$.

Thus, to prove that $f - g < f - h$, it suffices to prove that $\beta^i_\delta(f_m) - \beta^i_\delta(f_m - \rho_m) \le \beta^i_\delta(f_{m+1}) - \beta^i_\delta(f_{m+1} - \rho_m)$. Note that $f_m = \beta_\delta(f_{m+1}) = f_{m+1} + \mathit{gain}_\delta(f_{m+1}) > f_{m+1}$, because $\mathit{gain}_\delta(f_{m+1}) > 0$ by Claim 4 (since we are assuming $f_{m+1} > \alpha_g$, and $\alpha_g > \alpha^{\mathsf{min}}_\delta$). Note also that β^i_δ, as a self-composition of a concave increasing function, is concave by repeated application of Claim 17. Since a concave function is more sensitive to a change ρ_m in the input when the input is smaller, the result follows. Formally, the result follows by Claim 18 applied to $x_1 = f_{m+1}$, $x_2 = f_m$, and $z = \rho_m$. Because the results on concave functions are standard, general, and separate from the rest of the proof, we present them in Appendix B.

Case 2: 0s among footprints. Now we will deal with possible 0s among the f_i, g_i, and h_i values. Recall that if any of these values becomes 0 at some level, then it remains 0 at higher levels (so, conversely, if it is nonzero at some level, it is also nonzero below). We already are considering only the case when $f_{m+1} > \alpha_g$, so $f_{m+1} > \alpha^{\mathsf{min}}_\delta$, and thus we know by applying Claim 4 that $f_1 > \cdots > f_m > f_{m+1}$ (because there are no black pebbles on levels $1, \ldots, m$), so none of the f_i values is 0.

Claim 7. *For all i with $1 \le i \le m$, $g_i \ge h_{i+1}$.*

Proof. We will proceed by induction starting at $i = m$ and going down to $i = 1$. For the base case, note that $f_m = \beta_\delta(f_{m+1}) = f_{m+1} + \mathit{gain}_\delta(f_{m+1}) > f_{m+1}$, because we are considering only the case when $f_{m+1} > \alpha_g$, so we can apply Claim 4. Therefore, $g_m = \max(0, f_m - \rho_m) \ge \max(0, f_{m+1} - \rho_m) = h_{m+1}$.

The inductive step follow by monotonicity of β_δ, because for $i < m$, $g_i = \max(0, \beta_\delta(g_{i+1}))$ and $h_{i+1} = \max(0, \beta_\delta(h_{i+2}))$. □

Applying this claim, $g = \sum_{i=1}^{b} g_i = \sum_{i=m+2}^{b} g_i + g_{m+1} + \sum_{i=1}^{m} g_i \ge \sum_{i=m+2}^{b} h_i + 0 + \sum_{i=1}^{m} h_{i+1} = h - h_1$. If any of the g_i values is ever 0, then $g_1 = 0$, so by the above claim $h_2 = 0$ (since $h_2 \le g_1$), so $h_1 = 0$ and we are done. Similarly, if any of the h_i values is ever 0, then h_1 is 0 and we are done.

This concludes the proof of Lemma 3. □

7 Upperbounding the Number of Fertile Levels that Stop Growing

Recall from Step 3 in Section 3.1 that a fertile level starts with a footprint of at least α_π. Thanks to Theorem 3, we know how the footprint of a fertile level grows. Unfortunately,

the adversary can stop the growth completely by spending enough pebbles at some level to cover up the entire footprint. Intuitively, doing so will reduce the number of available black pebbles, so the adversary cannot do so too many times. But this intuition, even if we could make it formal, is insufficient: if the adversary could, just once in the middle of the graph, stop the growth of all fertile levels below, then the best we could hope for is a footprint of size half the graph, while we are aiming for a footprint that is almost the entire graph.

A stronger intuitive statement is that stopping the growth of a fertile level becomes more expensive the longer you wait, and becomes impossible if you wait too long. Formalizing it requires defining what it means to "wait" and to "stop" the growth. Recalling that $f_b \geq \alpha_\pi$ (per Step 3 in Section 3.1), we will consider the growth stopped at level i if $f_i(f_b, b) < \alpha_\pi$. (While this will give a slightly suboptimal bound, because such a footprint may yet recover as we got up, we are only slightly undercounting the cost to the adversary: note that the f_{b-1} can be dropped to 0 with $\beta_\delta(\alpha_\pi)$ black pebbles, while to get f_{b-1} to be below α_π takes at least g_{α_π} black pebbles, and these values are close for small α_π.) We thus provide the following definition.

Definition 5. We will say that level b is *viable* for k levels if $f_{b-i}(\alpha_\pi, b) \geq \alpha_\pi$ for all $0 \leq i < k$. If $m \geq b - k$, we will say that m is a *viable ancestor* of b. We will say that level b is *extinguished* after k levels if it is viable for k levels and $f_{b-k}(\alpha_\pi, b) < \alpha_\pi$.

The main idea for avoiding a messy case analysis based on different adversarial strategies is to think not about per-level footprint size and pebble allocation, but rather about the gap between the footprint and the number of available pebbles, thus reducing the problem to a single variable. The main result of this section is the following theorem, which uses this idea to show that the footprint becomes big after just a constant number of fertile levels.

Theorem 4. *Assume $g_{\alpha_\pi} > 0$. Let m be a fertile level; assume there are k fertile levels up to and including m. Fix some σ so that $\rho + \sigma < \alpha_\delta^{\mathsf{max}}$ and $\sigma > \alpha_\delta^{\mathsf{min}}$ (to understand this theorem, it may help to think of $\sigma = \varepsilon_{\mathsf{space}}/2$; thus, $\rho + \sigma = 1 - \varepsilon_{\mathsf{space}}/2$). Assume*

$$k \geq \max\left(\frac{\rho + \sigma - \alpha_\pi}{g_{\alpha_\pi}}, \beta\mathrm{count}_{\alpha_\pi}(\rho + \sigma)\right).$$

Then there is a fertile level $b \geq m$ for such that every unpebbled subset of level b of weight α_π has footprint at least $\alpha_\pi + \sum_{i=0}^{m-2} \beta_\delta^i(\sigma)$.

We defer the proof of this theorem to the end of the section. Note that m may be a viable ancestor of several different levels b; this theorem does not tell us which one we can choose in such a situation—it tells us only that one of them will work.

The central technical piece of the proof of Theorem 4 is the following lemma about the cost of viable levels. The main insight is to focus on the gain, rather than the footprint or black pebbles weight of each level. Intuitively, this approach works because the gain is easier to bound, and because to slow down the growth of the footprint (perhaps even to stop it completely), the adversary has to overcome the total gain, by Claim 3, no matter how the pebbles are allocated among levels.

Lemma 4. *Assume $g_{\alpha_\pi} \geq 0$. Let $f_b = \alpha_\pi$. Assume level b is viable for k levels. Then the total of first k gains satisfies*

$$gain_\delta(f_b) + \cdots + gain_\delta(f_{b-(k-1)}) \geq \min(k \cdot g_{\alpha_\pi}, \beta^k_\delta(\alpha_\pi) - \alpha_\pi)\,.$$

Interpretation of Lemma 4. Note that the sum of the first k gains is a function of $k-1$ black pebble weights $\rho_{b-1} \ldots \rho_{b-(k-1)}$. This lemma says that the minimum of this function, subject to the viability constraint, is at one of two extremal points of its domain: when $\rho_{b-1} = \rho_{b-2} = \cdots = \rho_{b-(k-1)} = g_{\alpha_\pi}$, or when $\rho_{b-1} = \cdots = \rho_{b-(k-1)} = 0$. In other words, if the adversary's goal is to minimize the gain while maintaining viability, the adversary can accomplish this goal by either spending enough black pebbles at each level to bring f_i value down to α_π for each i, or no black pebbles at all, to let f_i grow as fast as possible. Note that the bound given by this lemma is tight.

Proof of Lemma 4. Suppose for every m such that $b-k < m \leq b$, we have $gain_\delta(f_m) \geq g_{\alpha_\pi}$. Then we are done because the total gain for k levels is at least $k \cdot g_{\alpha_\pi}$.

Thus, the remaining case to consider is when for some m, $gain_\delta(f_m) < g_{\alpha_\pi}$. The following simple claim will be helpful.

Claim 8. *If for some i, $f_i \geq \alpha_g$, and there are no black pebbles above level i, then $gain_\delta(f_i) > gain_\delta(f_{i-1}) > \cdots > gain_\delta(f_1)$.*

Proof. Because there are no black pebbles, by Claim 4, $\alpha_g \leq f_i < f_{i-1} < \cdots < f_1$, and $gain_\delta$ is a decreasing function above α_g by Fact 1. □

We will now show a sequence of changes to the allocation of black pebbles. This sequence will be carefully constructed, so that each step in the sequence does not increase the total gain. At the end, the total gain will be at least as big as in the statement of the lemma.

1. Let m (with $b-k \leq m < b$) be the lowest level between b and $b-k$ with $gain_\delta(f_m) < g_{\alpha_\pi}$. Observe that this means $f_m > \alpha_g$ (by Fact 1, because $f_m \geq \alpha_\pi$ by viability, but $gain_\delta(f_m) < g_{\alpha_\pi}$).

   If there are any black pebbles at level m and above, remove them. Doing so will not decrease any of $f_{m-1}, \ldots, f_{b-(k-1)}$ (by Claim 2). Moreover, each of these f_i values will become greater that f_m by Claim 4 (because $f_m > \alpha_g > \alpha^{\mathsf{min}}_\delta$) and therefore also greater than α_g. Thus, if before this change, f_i was above α_g, then increasing f_i decreases its gain by Fact 1. Else, f_i was between α_π and α_g, and therefore $gain_\delta(f_i)$ was at least g_{α_π} by Fact 1, and it becomes smaller than $gain_\delta(f_m) < g_{\alpha_\pi}$ by Fact 1 and therefore decreases.

   Note the importance of removing all black pebbles at m and above at once: removing black pebbles one level at a time (either from $b-(k-1)$ down to m or from m to $b-(k-1)$) would not allow this argument to go through, as some f_i values may increase but not go above α_g.

2. Now proceed removing all black pebbles one level at a time from level $m-1$ down to $b-1$, in order, as long as removing all black pebbles at that level does not increase the total gain. If we get to level $b-1$, we are done, because the total gain is $f_{b-k} - f_b$ by Claim 3, which is $\beta_\delta^k(f_b) - f_b$ because there are no black pebbles. Else, let j be the the level at which this process stops: setting $\rho_j = 0$ increases the total gain, even though there are no longer any black pebbles above level j.

3. Consider the total gain of levels j through $b-(k-1)$: $\sum_{i=b-(k-1)}^{j} gain_\delta(f_i) = \beta_\delta^{j-(b-k)}(f_j) - f_j$. Consider this total gain as a function of f_j, where all the black pebble weights are fixed, except ρ_j. Note that $\beta_\delta^{j-(b-k)}$ and $-f_j$ are both concave functions of f_j (the former is by Claim 17, the latter because it's a line), and thus their sum is concave by 16, and thus the minimum is reached at the extrema of f_j by Claim 15. The largest f_j happens when $\rho_j = 0$, but we know, by the previous step, that removing all pebbles at level j increases the total gain, so $\rho_j = 0$ cannot give the minimum total gain. The smallest f_j is α_π, by viability, and thus the minimum possible total gain happens when $f_j = \alpha_\pi$. Note, again, the importance of the careful ordering of steps: we are using the fact that there are no black pebbles above level j, which implies (by Claim 4, which applies because $f_j > \alpha_\pi > \alpha_\delta^{\mathsf{min}}$) that $f_j < f_{j-1} < \cdots < f_{b-(k-1)}$, and thus as long as viability holds at level j, it also holds above up to level $b-(k-1)$; without the removal of pebbles above level j, the minimum allowed f_j could be larger than α_π due to viability constraints on the levels above.

   Thus, setting $\rho_j = \beta_\delta(f_{j+1}) - \alpha_\pi$ so that $f_j = \alpha_\pi$ will not increase the total gain. We do so. The total gain is now equal to $\sum_{i=j-1}^{b} gain_\delta(f_i) + \beta_\delta^{j-(b-k)}(\alpha_\pi) - \alpha_\pi$.

4. By the choice of m in the first step, we know $gain_\delta(f_i) \geq g_{\alpha_\pi}$ for $j < i < b$ (because black pebble quantities at levels below j have not been changed yet). Note that this step crucially uses that m was chosen as the *lowest* level with $gain_\delta(f_m) < g_{\alpha_\pi}$. By the step above, $gain_\delta(f_j) = g_{\alpha_\pi}$. Above f_j, the f values are increasing (by Claim 4). If $gain_\delta$ above f_j is always at least g_{α_π}, the total gain is at least $k \cdot g_{\alpha_\pi}$ and we are done.

   Else for some level $j' < j$, $gain_\delta(f_{j'}) < g_{\alpha_\pi}$, which means $f_{j'} > \alpha_g$ (because $f_{j'} \geq \alpha_\pi$ by viability, so if $f_{j'} \leq \alpha_g$, then $gain_\delta(f_{j'}) \geq g_{\alpha_\pi} = g_{\alpha_\pi}$ by Fact 1).

   In such a case, remove all the remaining black pebbles above level b. This shifts the f values down by $b-j$ steps. That is, the new values of $f_b, \ldots, f_{b-(k-1)+(b-j)}$ become equal to the old values of $f_j \ldots f_{b-(k-1)}$ (since $f_b = \alpha_\pi$ and now there are no black pebbles above b, just like before the removal of the pebbles, f_j was α_π and there were no black pebbles above j). The new $b-j$ values $f_{(b-k)+(b-j)}, \ldots, f_{b-(k-1)}$ all have gains less than g_{α_π} by the existence of j' and Claim 8), whereas before this removal of pebbles, the $b-j$ values $f_b, \ldots, f_{j+1}$ had gains greater than g_{α_π}. Thus, the total gain does not increase, because we removed $b-j$ levels at the bottom whose gain was at least g_{α_π}, shifted $k-(b-j)$ levels down without changing the gains, and added $b-j$ levels at the top whose gain is less than g_{α_π}. But, by Claim 3, the total gain is now $f_{b-k} - f_b = \beta_\delta^k(\alpha_\pi) - \alpha_\pi$.

This concludes the proof of Lemma 4. □

This lemma tells us, in particular, what it takes to extinguish a viable level.

Corollary 1. *Assume $g_{\alpha_\pi} \geq 0$. Assume level b is extinguished after k levels. Then*

$$\rho_{b-1\ldots b-k} \geq \min(k \cdot g_{\alpha_\pi}, \beta_\delta^k(\alpha_\pi) - \alpha_\pi)\,.$$

Proof. By Claim 3, $f_{b+k} \geq \alpha_\pi + \sum_{i=b-(k-1)}^{b} gain_\delta(f_i) - \rho_{b-1\ldots b-k}$. Since b is extinguished, $\alpha_\pi > f_{b+k}$, so

$$\rho_{b-1\ldots b-k} > \sum_{i=b-(k-1)}^{b} gain_\delta(f_i) \geq \min(k \cdot g_{\alpha_\pi}, \beta_\delta^k(\alpha_\pi) - \alpha_\pi)$$

by Lemma 4. □

The following corollary, in contrast to Corollary 1, speaks of levels that have not been extinguished. We cannot bound the number of pebbles spent on such levels, but we can bound the sum of the number of pebbles and the expansion of the last level.

Corollary 2. *Assume $g_{\alpha_\pi} \geq 0$. Assume level b is viable for k levels, and $m = b-(k-1)$. Then*

$$\rho_{b-1\ldots m} + \beta_\delta(f_m) \geq \alpha_\pi + \min(k \cdot g_{\alpha_\pi}, \beta_\delta^k(\alpha_\pi) - \alpha_\pi)\,.$$

Proof. By Claim 3 $\beta_\delta(f_m) = f_m + gain_\delta(f_m) \geq \alpha_\pi + \sum_{i=m}^{b} gain_\delta(f_i) - \rho_{b-1\ldots m}$, so

$$\rho_{b-1\ldots m} + \beta_\delta(f_m) \geq \alpha_\pi + \sum_{i=m}^{b} gain_\delta(f_i) \geq \alpha_\pi + \min(k \cdot g_{\alpha_\pi}, \beta_\delta^k(\alpha_\pi) - \alpha_\pi)$$

by Lemma 4. □

Consider now several levels that are extinguished after some number of levels each. Add up the total black pebbles required. The following (rather boring and technical) claim shows that what you get is at least as big as if you had a single level extinguished after the combined total number of levels.

Claim 9. *Assume $g_{\alpha_\pi} \geq 0$. Let k_1 and k_2 be positive integers. Then*

$$\begin{aligned}
&\min((k_1 + k_2) \cdot g_{\alpha_\pi}, \beta_\delta^{k_1+k_2}(\alpha_\pi) - \alpha_\pi) \\
\leq &\min(k_1 \cdot g_{\alpha_\pi}, \beta_\delta^{k_1}(\alpha_\pi) - \alpha_\pi) \\
&+ \min(k_2 \cdot g_{\alpha_\pi}, \beta_\delta^{k_2}(\alpha_\pi) - \alpha_\pi)\,.
\end{aligned}$$

Proof. Intuitively, as chains get longer, per level gains eventually start decreasing, and longer chains have more time to benefit from this decrease. Now we give the formal proof.

If $k_1 \cdot g_{\alpha_\pi} \leq \beta_\delta^{k_1}(\alpha_\pi) - \alpha_\pi$ and $k_2 \cdot g_{\alpha_\pi} \leq \beta_\delta^{k_2}(\alpha_\pi) - \alpha_\pi$, then the sum in question is equal to $(k_1 + k_2) \cdot g_{\alpha_\pi}$ and we are done.

Else, assume, without loss of generality, that $\beta_\delta^{k_1}(\alpha_\pi) - \alpha_\pi < k_1 \cdot g_{\alpha_\pi}$. Note that $\beta_\delta^{k_1}(\alpha_\pi) - \alpha_\pi = \sum_{i=0}^{k_1-1} gain_\delta(\beta_\delta^i(\alpha_\pi))$ by definition of $gain_\delta$. Therefore, for some m (with $0 \le m < k_1$), $gain_\delta(\beta_\delta^m(\alpha_\pi)) < g_{\alpha_\pi}$, which, by Fact 1, means $\beta_\delta^m(\alpha_\pi) > \alpha_g$ (because $\beta_\delta^m(\alpha_\pi) \ge \alpha_\pi$ by Claim 4). Take the smallest such m. By Claim 8, $gain_\delta(\beta_\delta^{j_1}(\alpha_\pi)) \le gain_\delta(\beta_\delta^{j_2}(\alpha_\pi)) < g_{\alpha_\pi}$ for any $j_1 \ge j_2 \ge m$. From this step, we derive two inequalities.

- Because $k_1 \ge m$, $\sum_{i=k_1}^{k_1+k_2-1} gain_\delta(\beta_\delta^i(\alpha_\pi)) < k_2 \cdot g_{\alpha_\pi}$. Therefore,

$$\begin{aligned}\beta_\delta^{k_1+k_2}(\alpha_\pi) - \alpha_\pi &= \sum_{i=0}^{k_1-1} gain_\delta(\beta_\delta^i(\alpha_\pi)) + \sum_{i=k_1}^{k_1+k_2-1} gain_\delta(\beta_\delta^i(\alpha_\pi)) \\ &\le \sum_{i=0}^{k_1-1} gain_\delta(\beta_\delta^i(\alpha_\pi)) + k_2 \cdot g_{\alpha_\pi} \\ &= (\beta_\delta^{k_1}(\alpha_\pi) - \alpha_\pi) + k_2 \cdot g_{\alpha_\pi}\,.\end{aligned}$$

- Take any $i \ge 0$. Set $j_1 = i + k_1$ and $j_2 = i$. Note that $j_1 \ge m$ because $k_1 \ge m$. We can show by cases that $gain_\delta(\beta_\delta^{i+k_1}(\alpha_\pi)) \le gain_\delta(\beta_\delta^i(\alpha_\pi))$, as follows: if $i = j_2 \ge m$, we have already shown it, and if $i = j_2 < m$, then $gain_\delta(\beta_\delta^{j_2}(\alpha_\pi)) \ge g_{\alpha_\pi}$, while $gain_\delta(\beta_\delta^{j_1}(\alpha_\pi)) \ge g_{\alpha_\pi}$ because $j_1 \ge m$. Therefore,

$$\begin{aligned}\beta_\delta^{k_1+k_2}(\alpha_\pi) - \alpha_\pi &= \sum_{i=0}^{k_1-1} gain_\delta(\beta_\delta^i(\alpha_\pi)) + \sum_{i=k_1}^{k_1+k_2-1} gain_\delta(\beta_\delta^i(\alpha_\pi)) \\ &\le \sum_{i=0}^{k_1-1} gain_\delta(\beta_\delta^i(\alpha_\pi)) + \sum_{i=0}^{k_2-1} gain_\delta(\beta_\delta^i(\alpha_\pi)) \\ &= (\beta_\delta^{k_1}(\alpha_\pi) - \alpha_\pi) + (\beta_\delta^{k_2}(\alpha_\pi) - \alpha_\pi)\,.\end{aligned}$$

These two inequalities together conclude the proof of Claim 9. □

Proof of Theorem 4. The assumptions imply that the βcount term is finite by Claim 1, because $g_{\alpha_\pi} \ge 0$ implies $\alpha_\pi > \alpha_\delta^{\mathsf{min}}$.

Starting with level ℓ and going up, find the lowest fertile level b_1; assume level b_1 becomes extinguished after k_1 levels, and $b_1 - k_1 \ge m$. This gives us a lower bound

$$\rho_{b_1-1\ldots b_1-k_1} \ge \min(k_1 \cdot g_{\alpha_\pi}, \beta_\delta^{k_1}(\alpha_\pi) - \alpha_\pi)\,,$$

by Corollary 1. Skip infertile levels (if any) at or above $b_1 - k_1$ to find a fertile level $b_2 \le b_1 - k_1$, and assume level b_2 becomes extinguished after k_2 levels, and $b_2 - k_2 \ge m$. This, again, gives us a lower bound

$$\rho_{b_2-1\ldots b_2-k_2} \ge \min(k_2 \cdot g_{\alpha_\pi}, \beta_\delta^{k_2}(\alpha_\pi) - \alpha_\pi)\,,$$

Note that the regions for which we obtain these bounds on black pebble weight do not overlap, as $b_2 - 1 < b_1 - k_1$. Note also that we will not skip over m, as it is fertile.

Continuing in this manner, eventually we will come to a fertile level $b \geq m$ that stays viable until level m inclusive. Then, letting $k' = b - m + 1$

$$\rho_{b-1\ldots m} + \beta_\delta(f_m(\alpha_\pi, b)) \geq \alpha_\pi + \min(k' \cdot g_{\alpha_\pi}, \beta_\delta^{k'}(\alpha_\pi) - \alpha_\pi)$$

by Corollary 2.

Adding up all the inequalities per Claim 9 and observing that the bounds on ρ values are for nonoverlapping ranges of levels, we obtain

$$\rho_{\ell\ldots m} + \beta_\delta(f_m(\alpha_\pi, b)) \geq \alpha_\pi + \min((k_1 + k_2 + \cdots + k') \cdot g_{\alpha_\pi}, \beta_\delta^{k_1+k_2+\cdots+k'}(\alpha_\pi) - \alpha_\pi)\,.$$

Note that $k_1 + k_2 + \cdots + k' \geq k$, because the only levels we skipped were infertile (we didn't necessarily skip all infertile levels, as some of them may have been viable; hence the inequality rather than equality). Replacing $k_1 + k_2 + \cdots + k'$ with k on the right-hand side will not increase it. Noting that $\rho_{\ell\ldots m} = \rho - \rho_{1\ldots m-1}$, we thus obtain

$$\beta_\delta(f_m(\alpha_\pi, b)) - \rho_{1\ldots m-1} \geq \min(\alpha_\pi + k \cdot g_{\alpha_\pi}, \beta_\delta^k(\alpha_\pi)) - \rho\,.$$

By the condition on k in Theorem 4, the right-hand side of this inequality at least σ. We can thus apply Theorem 3 to level m and substitute σ instead of $\beta_\delta(f_m(\alpha_\pi, b)) - \rho_{1\ldots m-1}$ by monotonicity of β (the condition $f_m \geq \alpha_\pi$ is satisfied because f_m is viable). □

8 Finishing the Proof with Quantitative Details for the Chung Expander

We explain and analyze Chung expanders in detail in Appendix A. Here we state only the fact that we need for the proof of Theorem 1.

Fact 2. *For any δ, a random degree-8 Chung expander of size $\tilde{O}(1/\delta)$ ensures, with overwhelming probability, that*

1. $\alpha_g \approx 0.32$
2. $gain(\alpha_g) \approx 0.36$
3. *Expansion is rapid for small sets: for any α satisfying $\delta \leq \alpha \leq 0.14$, $\beta(\alpha) > 3\alpha$, and therefore $\beta_\delta(\alpha) > 2\alpha$.*
4. *Expansion quickly bridges small-to-big gap: $\beta(\beta(\beta(0.14) - 0.14) - 0.14) - 0.14 > 0.68$, i.e., $\beta_\delta^3(0.14) > 0.68$ assuming $\delta \leq 0.14$*
5. *Expansion quickly reaches almost everything for big sets: for any α satisfying $\delta \leq \alpha \leq 0.14$, $\beta(1 - 3\alpha) > 1 - \alpha$. Therefore $\beta_\delta(1 - 3\alpha) > 1 - 2\alpha$, and moreover, $\beta_\delta(1 - 3\alpha) > 1 - 3\alpha/2$ as long as $2\delta \leq \alpha \leq 0.14$.*

(Note that the last item can be inferred from item 3 via Claim 13: we have $\beta(\alpha) \geq 3\alpha$, so, by monotonicity of β, $\beta(1 - 3\alpha) > \beta(1 - \beta(\alpha)) = 1 - \alpha$.)

These facts imply that for any $\delta \leq 0.14$, if there are no black pebbles, expansion can rapidly get the per layer footprint from δ to $1 - 3\delta$, as shown in the following claim.

Claim 10. *For the degree-8 Chung expander, $\delta \le 0.14$, and any $a \ge \delta$ and $b \ge 3\delta$*

$$\beta\mathrm{count}_a(1-b) \le 3 + \lceil \log_2(0.14/a) \rceil + \lceil \log_2(0.32/b) \rceil < 1 + \log_2 1/a + \log_2 1/b\,.$$

Proof. Let $k_1 = \lceil \log_2(0.14/a) \rceil$. By Fact 2.3, $\beta_\delta(\alpha) > 2\alpha$ for any α with $\delta \le \alpha \le 0.14$, so, by monotonicity of β (Condition 1), $\beta_\delta^{k_1}(a) > \min(2 \cdot 0.14, a \cdot 2^{k_1}) \ge \min(0.28, a \cdot 2^{\log_2(0.14/a)}) = 0.14$. By Fact 2.4 and monotonicity of β, $\beta_\delta^{k_1+3}(a) > 0.68 = 1 - 0.32$.

Let $k_2 = \lceil \log_2(0.32/b) \rceil$. By Fact 2.5, $\beta_\delta^{k_1+k_2+3}(a) = \beta_\delta^{k_2}(\beta_\delta^{k_1+3}(a)) > \beta_\delta^{k_2}(1-0.32) > \min(1 - 3 \cdot \delta, 1 - 0.32/2^{k_2}) > 1 - b$. □

We are now ready to prove Theorem 1, which we restate here.

Theorem 1. *For any $\varepsilon_{\mathsf{space}} > 0$ and $\varepsilon_{\mathsf{hardness}} > 0$, there is a setting of parameters ℓ, ℓ_{pr}, δ, and n in the SPR construction that achieves space gap $\varepsilon_{\mathsf{space}}$, single-query catching probability $p_{\mathsf{hard}} \ge \varepsilon_{\mathsf{space}}/2$, and hardness gap $\varepsilon_{\mathsf{hardness}}$, such that*

- *The cost of computing $f(x)$, per bit of y, is*

$$O\left(\frac{1}{\varepsilon_{\mathsf{hardness}}} \cdot \log \frac{1}{\varepsilon_{\mathsf{space}}}\right).$$

- *The cost of the initialization protocol is*

$$\tilde{O}\left(\frac{1}{\varepsilon_{\mathsf{hardness}}} \cdot \left(\frac{1}{\varepsilon_{\mathsf{hardness}}} + \frac{1}{\varepsilon_{\mathsf{space}}}\right)\right).$$

.

Proof. Set $\zeta = 1 - \varepsilon_{\mathsf{space}}/2 = \rho + \varepsilon_{\mathsf{space}}/2$ and $\sigma = \varepsilon_{\mathsf{space}}/2$. Set δ to satisfy the following conditions:

$$\delta < \min\left(\frac{gain(\pi)}{2}, \frac{\varepsilon_{\mathsf{space}}}{6}, \frac{1-\pi}{3}, \frac{gain(\alpha_\pi)}{2}, \alpha_\pi, \frac{\varepsilon_{\mathsf{hardness}}}{6}, 0.14\right)$$

Set m_1 to be the largest integer smaller than

$$\max\left(1 + \frac{\pi + \delta - \varepsilon_{\mathsf{space}}/2}{g_\pi}, 1 + \beta\mathrm{count}_{\varepsilon_{\mathsf{space}}/2-\delta}(\pi)\right).$$

The first condition on δ (namely, $\delta < \frac{gain(\pi)}{2}$) ensures that the denominator $g_\pi = gain_\delta(\pi) = gain(\pi) - \delta$ in the calculation of m_1 is at least $gain(\alpha_\pi)/2$ and thus constant (note that the numerator is at most 1). The second, third, and last condition ensure that $\beta\mathrm{count}_{\varepsilon_{\mathsf{space}}/2-\delta}(\pi)$ is logarithmic in $1/\varepsilon_{\mathsf{space}}$ by Claim 10, because $\varepsilon_{\mathsf{space}}/2 - \delta > \delta$ and $\pi < 1 - 3\delta$ is a constant. Thus, m_1 is most logarithmic in $1/\varepsilon_{\mathsf{space}}$.

Set m_2 to be the largest integer no greater than

$$\max\left(\frac{\rho + \sigma - \alpha_\pi}{gain_\delta(\alpha_\pi)}, \beta\mathrm{count}_{\alpha_\pi}(\rho + \sigma)\right).$$

A similar argument, using the fourth, fifth, and second condition on δ, and remembering that $\rho + \sigma = 1 - \varepsilon_{\mathsf{space}} + \varepsilon_{\mathsf{space}}/2 = 1 - \varepsilon_{\mathsf{space}}/2 < 1 - 3\delta$, shows that m_2 is logarithmic in $1/\varepsilon_{\mathsf{space}}$.

Set $\ell_{\mathsf{pr}} = m_1 + m_2$. Somewhere among the lowest ℓ_{pr} levels, there must be a fertile level with $m_2 - 1$ fertile levels below it, because if not, then the total number of fertile levels among the lowest ℓ_{pr} levels is less than m_2, so the total number of infertile levels is greater than m_1, which contradicts Theorem 2.

Therefore, Theorem 4 applies (because $\rho + \sigma < 1 - \delta < \alpha_\delta^{\mathsf{max}}$ and $\sigma > \delta > \alpha_\delta^{\mathsf{min}}$), which means that among the lowest ℓ_{pr} levels, there is a fertile level b so that every unpebbled set of weight α_π of level b has footprint at least

$$\alpha_\pi + \sum_{i=0}^{\ell - \ell_{\mathsf{pr}} - 1} \beta_\delta^i(\sigma)\,. \tag{2}$$

It remains to find a lower bound on (2). Let $m_3 = \beta\mathrm{count}_\sigma(0.68)$ (this is again logarithmic in $1/\varepsilon_{\mathsf{space}}$ by Claim 10). Ignore the first m_3 terms of the summation in (2). Now let $m_4 = \lceil \log_2 0.32/(\varepsilon_{\mathsf{hardness}}/2) \rceil$. The next m_4 terms of the summation add up to at least

$$(1 - 0.32) + (1 - 0.32/2) + (1 - 0.32/2^2)) + \cdots + (1 - 0.32/2^{m_4 - 1}) > m_4 - 1$$

by Fact 2.5, because $3\delta < \varepsilon_{\mathsf{hardness}}/2$ by the sixth condition on δ. Each subsequent term is at least $1 - \varepsilon_{\mathsf{hardness}}/2$, and so the lower bound on (2) is

$$(\ell - \ell_{\mathsf{pr}} - m_3 - 1) \cdot (1 - \frac{\varepsilon_{\mathsf{hardness}}}{2})\,.$$

Setting $\ell = \frac{2(m_1 + m_2 + m_3 + 1)}{\varepsilon_{\mathsf{hardness}}} = O(\frac{1}{\varepsilon_{\mathsf{hardness}}} \log \frac{1}{\varepsilon_{\mathsf{space}}})$ gives that this footprint, relative to the weight ℓ of the entire graph, is greater than $1 - \varepsilon_{\mathsf{space}}$. The cost of computing $f(x)$ per output byte of y is proportional to ℓ, and is thus

$$O\left(\frac{1}{\varepsilon_{\mathsf{hardness}}} \cdot \log \frac{1}{\varepsilon_{\mathsf{space}}}\right)\,.$$

Per Fact 2, n grows as $\tilde{O}(1/\delta)$ in order to ensure expansion works as desired for $\alpha_\delta^{\mathsf{min}} < \alpha < \alpha_\delta^{\mathsf{max}}$ (though for practically relevant parameter values, a typical n dictated by the desired storage size will be sufficient for any reasonably relevant δ); per [ABP18], the degree of depth-robust graphs grows logarithmically in n (it may be that predecessor-robust graphs can have a smaller dependence on n than depth-robust graphs, but we do not know that). Thus, d_{pr} grows at most logarithmically in $1/\delta$; d_{exp} is the constant 8.

Therefore, the intialization complexity, per Section 2.2 is

$$O\left(\frac{\ell_{\mathsf{pr}} \cdot d_{\mathsf{pr}} + \ell \cdot d_{\mathsf{exp}}}{\delta}\right) = \tilde{O}\left(\frac{1}{\varepsilon_{\mathsf{hardness}}} \cdot \left(\frac{1}{\varepsilon_{\mathsf{hardness}}} + \frac{1}{\varepsilon_{\mathsf{space}}}\right)\right)\,.$$

□

Acknowledgments

Thanks to Carla Rafòls and Anca Nitulescu for being patient sounding boards for my half-baked ideas, to Irene Giacomelli and Luca Nizzardo for helping me understand the existing SDR proof, and to Matteo Campanelli, Nicola Greco, and Ben Fisch for helping to connect me with this problem. Anonymous referees provided helpful suggestions for improving the manuscript. Thanks to Protocol Labs for funding this work and to Carla Rafòls for hosting me at Universitat Pompeu Fabra.

References

[AAC+17] Hamza Abusalah, Joël Alwen, Bram Cohen, Danylo Khilko, Krzysztof Pietrzak, and Leonid Reyzin. Beyond hellman's time-memory trade-offs with applications to proofs of space. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part II*, volume 10625 of *LNCS*, pages 357–379. Springer, Cham, December 2017.

[ABFG14] Giuseppe Ateniese, Ilario Bonacina, Antonio Faonio, and Nicola Galesi. Proofs of space: When space is of the essence. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 538–557. Springer, Cham, September 2014.

[ABP18] Joël Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. Sustained space complexity. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 99–130. Springer, Cham, April / May 2018.

[AdNV17] Joël Alwen, Susanna F. de Rezende, Jakob Nordström, and Marc Vinyals. Cumulative space in black-white pebbling and resolution. In Christos H. Papadimitriou, editor, *ITCS 2017*, volume 4266, pages 38:1–38:21, 67, January 2017. LIPIcs.

[AGK+18] Joël Alwen, Peter Gazi, Chethan Kamath, Karen Klein, Georg Osang, Krzysztof Pietrzak, Leonid Reyzin, Michal Rolinek, and Michal Rybár. On the memory-hardness of data-independent password-hashing functions. In Jong Kim, Gail-Joon Ahn, Seungjoo Kim, Yongdae Kim, Javier López, and Taesoo Kim, editors, *ASIACCS 18*, pages 51–65. ACM Press, April 2018.

[BDG17] Juan Benet, David Dalrymple, and Nicola Greco. Proof of replication, 2017. `https://filecoin.io/proof-of-replication.pdf`.

[Chu79] F.R.K. Chung. On concentrators, superconcentrators, generalizers, and nonblocking networks. *Bell System Technical Journal*, 58(8):1765–1777, 1979.

[CP19] Bram Cohen and Krzysztof Pietrzak. The Chia network blockchain, 2019. Unpublished Manuscript.

[DFKP15] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 585–605. Springer, Berlin, Heidelberg, August 2015.

[EGS75] P. Erdős, R.L. Graham, and E. Szemerédi. On sparse graphs with dense long paths. Technical Report STAN-CS-75-504, Stanford University Computer Science Department, September 1975. `http://i.stanford.edu/pub/cstr/reports/cs/tr/75/504/CS-TR-75-504.pdf`.

[FBGB18] Ben Fisch, Josep Bonneau, Nicola Greco, and Juan Benet. Scaling proof-of-replication for filecoin mining, 2018. `https://research.protocol.ai/publications/scaling-proof-of-replication-for-filecoin-mining/fisch2018.pdf`.

[Fis18] Ben Fisch. PoReps: Proofs of space on useful data. Cryptology ePrint Archive, Report 2018/678, 2018.

[Fis19] Ben Fisch. Tight proofs of space and replication. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 324–348. Springer, Cham, May 2019.

[GN23] Irene Giacomelli and Luca Nizzardo. Filecoin proof of useful space — technical report. Technical report, CryptoNet — Protocol Labs, 2023. `https://drive.google.com/file/d/1notObdkPT1BCztgspIpzSUAzWSrM8h81/view`.

[Pie19] Krzysztof Pietrzak. Proofs of catalytic space. In Avrim Blum, editor, *ITCS 2019*, volume 124, pages 59:1–59:25. LIPIcs, January 2019.

[PPK+15] Sunoo Park, Krzysztof Pietrzak, Albert Kwon, Joël Alwen, Georg Fuchsbauer, and Peter Gaži. Spacemint: A cryptocurrency based on proofs of space. Cryptology ePrint Archive, Report 2015/528, 2015.

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

[Pro23] Protocol Labs. Filecoin spec. Algorithms. Stacked DRG PoRep, 2023. `https://spec.filecoin.io/algorithms/sdr/`.

[RD16] Ling Ren and Srinivas Devadas. Proof of space from stacked expanders. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 262–285. Springer, Berlin, Heidelberg, October / November 2016.

[Stă01] Pantelimon Stănică. Good lower and upper bounds on binomial coefficients. *Journal of Inequalities in Pure and Applied Mathematics*, 2, 2001. `https://www.emis.de/journals/JIPAM/images/043_00_JIPAM/043_00.pdf`.

A The function β for the Chung Expander

The Chung bipartite expander [Chu79] of degree d is a randomized construction, built as follows. Each of the two parts has n vertices; think of each vertex as having d ports; number all the ports sequentially (i.e., port i belongs to vertex $\lfloor i/d \rfloor$), choose a random permutation p, and draw an edge from port i on the top part to port $p(i)$ on the bottom part. In this section we describe the expansion function β and argue that a random Chung graph has this expansion function with overwhelming probability as n grows. The main results of this section are Equation (4) and Claim 14 below.

Chung expanders were first in the context of proofs of space by [RD16]. However, the work of [RD16] used expansion at one point α, whereas we want it to work almost the entire $[0, 1]$ — specifically, on the interval $[\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}]$. We will do so by taking the union bound over (at most n) possible values of α (for integer αn), which requires us to understand the probability that expansion fails more precisely than in prior work. Thus, we first reprove [RD16, Theorem 1] with more precise constants.

Claim 11. *For every fixed u and v, the probability (over the choice of the permutation p) that there exists a set of u nodes at the bottom that does not have at least v predecessors at the top is at most*

$$\frac{c}{n} \cdot 2^{n \cdot (\mathsf{H_b}(x) + \mathsf{H_b}(y) + d \cdot (y\mathsf{H_b}(x/y) - \mathsf{H_b}(x)))},$$

where $x = u/n$, $y = v/n$, $\mathsf{H_b}$ is the binary entropy function $\mathsf{H_b}(x) = -(x \log_2 x + (1 - x) \log_2(1 - x))$, and

$$c < \frac{\exp(1/8)}{2\pi} \cdot \frac{1}{\sqrt{x(1-y)(y-x)}}.$$

Proof. Before we proceed, it helps to introduce notation $C(a, b) = a!/b!/(b - a)!$ for combinations and $P(a, b) = a!/(b - a)!$ for permutations. We show a very tight bound on the binomial coefficient in the following claim (the slack in the bound is just 13%, because $\exp(-1/8) \approx 0.88$).

Claim 12. *For any $q \geq 1$ and $1/q \leq a < 1$, there is a value κ such that $\exp(-1/8) < \kappa < 1$ and*

$$C(q, aq) = \kappa \cdot \frac{1}{\sqrt{2\pi a(1-a)}} \cdot \frac{1}{\sqrt{q}} \cdot 2^{q\mathsf{H_b}(a)}.$$

Proof. The result of Stănică [Stă01, Theorem 2.6] tells us that for all $m > p \geq 1$ and for all $k \geq 1$,

$$C(mk, pk) = \kappa \cdot \sqrt{\frac{m}{2\pi k(m-p)p}} \left(\frac{m^m}{(m-p)^{m-p} \cdot p^p} \right)^k$$

for some value κ with $\exp(-1/(8k)) < \kappa < 1$.

Plugging in $p = 1, k = aq$, and $m = 1/a$ (note that so that $m > p \geq 1$ and $k \geq 1$ are satisfied), we have $\exp(-1/8) \leq \exp(-1/(8k)) < \kappa < 1$, $m - p = (1-a)/a$, and

$$\begin{aligned}
C(n, an) &= \kappa \cdot \sqrt{\frac{1/a}{2\pi aq(1-a)/a}} \left(\frac{(1/a)^{1/a}}{((1-a)/a)^{(1-a)/a}} \right)^{aq} \\
&= \kappa \cdot \sqrt{\frac{1}{2\pi qa(1-a)}} \left(\frac{(1/a)}{((1-a)/a)^{1-a}} \right)^{q} \\
&= \kappa \cdot \sqrt{\frac{1}{2\pi qa(1-a)}} \left(\frac{a^{-1} \cdot a^{1-a}}{(1-a)^{1-a}} \right)^{q} \\
&= \kappa \cdot \sqrt{\frac{1}{2\pi qa(1-a)}} \left(\frac{1}{a^{a}(1-a)^{1-a}} \right)^{q} \\
&= \kappa \cdot \sqrt{\frac{1}{2\pi qa(1-a)}} \cdot 2^{q\mathsf{H_b}(a)} .
\end{aligned}$$

This concludes the proof of Claim 12. □

We use the same reasoning as Ren and Devadas [RD16, Theorem 1]. Fix x and y. Call a bipartite graph (x, y)-*bad* if there exists a set of size xn in the bottom part that has fewer than yn predecessors in the top part. How many bad Chung graphs are there? To construct a bad Chung graph, there are at most $C(n, xn)$ candidates for the bottom set, at most $C(n, yn)$ candidates for the top set, at most $P(dyn, dxn)$ choices for where in the top set the edges that go to the bottom set originate, and $P(d(n - xn), d(n - xn)) = (dn - dxn)!$ choices for where the remaining edges originate. Since there are a total of $(dn)!$ possible graphs, the probability that a graph is (x, y)-bad is at most

$$\begin{aligned}
\Pr[\text{graph is } (x,y)\text{-bad}] &\leq C(n, xn) \cdot C(n, yn) \cdot P(dyn, dxn) \cdot (dn - dxn)!/(dn)! \\
&= C(n, xn) \cdot C(n, yn) \cdot C(dyn, dxn)/C(dn, dxn)
\end{aligned}$$

Thus, by Claim 12, as long as

$$1/n \leq x < 1, \quad 1/n \leq y < 1, \quad 1/(dyn) \leq x/y < 1, \text{ and } 1/(dn) \leq x < 1 \tag{3}$$

(note that the last two conditions are redundant, as they are both implied by the first), we have

$$\Pr[\text{graph is } (x,y)\text{-bad}] \leq c \cdot \frac{1}{n} \cdot 2^{n(\mathsf{H_b}(x) + \mathsf{H_b}(y) + d(y\mathsf{H_b}(x/y) - \mathsf{H_b}(x))}$$

for some c satisfying

$$\begin{aligned} c &< \exp(1/8)\sqrt{\frac{2\pi d \cdot x(1-x)}{(2\pi)^3 dy \cdot x(1-x) \cdot y(1-y) \cdot (x/y)(1-x/y)}} \\ &< \frac{\exp(1/8)}{2\pi}\sqrt{\frac{1}{y \cdot y(1-y) \cdot (x/y)(1-x/y)}} \\ &= \frac{\exp(1/8)}{2\pi}\sqrt{\frac{1}{x(1-y)(y-x)}}\,. \end{aligned}$$

This concludes the proof of Claim 11. □

Claim 11 tells us that if $\mathsf{H_b}(x) + \mathsf{H_b}(y) + d \cdot (y\mathsf{H_b}(x/y) - \mathsf{H_b}(x))$ is negative and n is sufficiently large, expansion is overwhelmingly likely.

We can thus define

$$\beta_{\mathsf{optmial}}(\alpha) = \sup\{y \,:\, \mathsf{H_b}(\alpha) + \mathsf{H_b}(y) + d \cdot (y\mathsf{H_b}(\alpha/y) - \mathsf{H_b}(\alpha)) < 0\}\,.$$

This function satisfies Condition 1 (we have verified this fact by plotting the function; analytical verification appears to be doable but painful and uninsightful). However, to get a handle on the probability that a random Chung graph of a given size n is an expander, we need a slightly stronger condition that simple negativity. We will have to pick a small value $\varepsilon_{\mathsf{chung}} << \delta$ (in our concrete example in Appendix C, $\varepsilon_{\mathsf{chung}} = 2^{-23}$) and define

$$\beta(\alpha) = \sup\{y \,:\, \mathsf{H_b}(\alpha) + \mathsf{H_b}(y) + d \cdot (y\mathsf{H_b}(\alpha/y) - \mathsf{H_b}(\alpha)) < -\varepsilon_{\mathsf{chung}}\}\,. \tag{4}$$

This change has a very small effect on β: $\beta_{\mathsf{optmial}}(\alpha) - \beta(\alpha) < \varepsilon_{\mathsf{chung}}$, because the derivative of the left-hand side of the inequality in (4) as a function of y is greater than 1 in the relevant region. As long as $\delta \gg \varepsilon_{\mathsf{chung}}$, this change is therefore not significant, because we are interested only in α for which $gain(\alpha) \geq \delta$.

The following claim helps understand expansion.

Claim 13. *$\beta(x) = y$ if and only if $\beta(1-y) = 1-x$*

Proof. $\mathsf{H_b}(1-y) = \mathsf{H_b}(y)$ and $\mathsf{H_b}(1-x) = \mathsf{H_b}(x)$ by definition of $\mathsf{H_b}$.

$$\begin{aligned} y\mathsf{H_b}(x/y) - \mathsf{H_b}(x) &= -y \cdot \left(\frac{x}{y} \cdot \log\frac{x}{y} + \frac{y-x}{y}\log\frac{y-x}{y}\right) \\ &\quad + x\log x + (1-x)\log(1-x) \\ &= -x\log x + x\log y - (y-x)\log(y-x) - x\log y + y\log y \\ &\quad + x\log x + (1-x)\log(1-x) \\ &= (x-y)\log(y-x) + y\log y + (1-x)\log(1-x)\,, \end{aligned}$$

which remains invariant if we substitute $1-x$ for y and $1-y$ for x. Thus, the inequality in (4) is true for (α, y) if and only if it is true for $(1-y, 1-\alpha)$. □

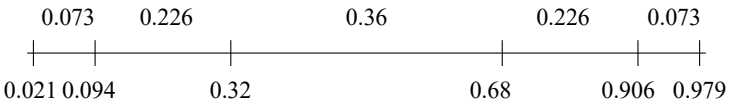

Figure 2: Expansion speed for the Chung expander: $\beta(0.021) \approx 0.094$, $\beta(0.094) \approx 0.32$, etc. The top numbers show the gain.

The above claim implies that for all α, $gain(\alpha) = \beta(\alpha) - \alpha = (1-\alpha) - (1-\beta(\alpha)) = \beta(1-\beta(\alpha)) - (1-\beta(\alpha)) = gain(1-\beta(\alpha))$. Therefore, $gain(1-\beta(\alpha_\delta^{\mathsf{min}})) = gain(\alpha_\delta^{\mathsf{min}}) = \delta$, so $\alpha_\delta^{\mathsf{max}} = 1 - \beta(\alpha_\delta^{\mathsf{min}})$ (because $\alpha_\delta^{\mathsf{max}}$ is the only value other than $\alpha_\delta^{\mathsf{min}}$ value with *gain* of δ).

The next claim shows that a random sufficiently large Chung graph is an expander with overwhelming probability.

Claim 14. *The probability that there exists $\alpha \in [\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}]$ and a set of nodes at the bottom of size $\alpha \cdot n$ whose predecessor set has size less than $\beta(\alpha) \cdot n$ is at most*

$$\frac{\exp(1/8)}{2\pi \cdot \alpha_\delta^{\mathsf{min}} \cdot \sqrt{\delta}} \cdot 2^{-n \cdot \varepsilon_{\mathsf{chung}}}.$$

Proof. Using Claim 11 and taking a union bound over all possible values of u (there are almost n of them, from $\alpha_\delta^{\mathsf{min}} \cdot n$ to $\alpha_\delta^{\mathsf{max}} \cdot n$), we get that the probability there exists an integer u and subset of weight $\alpha = u/n$ whose predecessor set weight is less than $\beta(\alpha)$ is at most $c_{\mathsf{max}} \cdot 2^{-n \cdot \varepsilon_{\mathsf{chung}}}$, where $c_{\mathsf{max}} = \sup_{x \in [\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}], y = \beta(x)} \frac{\exp(1/8)}{2\pi} \frac{1}{\sqrt{x(1-y)(y-x)}}$. Noting that $x \geq \alpha_\delta^{\mathsf{min}}$, $1 - y \geq 1 - \beta(\alpha_\delta^{\mathsf{max}}) = \alpha_\delta^{\mathsf{min}}$, and $y - x \geq \delta$, we have $c_{\mathsf{max}} \leq \frac{\exp(1/8)}{2\pi \alpha_\delta^{\mathsf{min}} \sqrt{\delta}}$. □

Thus, given a particular δ and a security parameter λ, we will set $\varepsilon_{\mathsf{chung}} \ll \delta$ to ensure Condition 1 holds on the interval $[\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}]$, and set

$$n > \left(\lambda - 2.4 - \log_2 \alpha_\delta^{\mathsf{min}} - \frac{1}{2} \log_2 \delta\right) / \varepsilon_{\mathsf{chung}} = \tilde{O}\left(\frac{1}{\delta}\right)$$

(sufficient because $\log_2(\exp(1/8)/(2\pi)) < -2.4$). This will ensure, with probability $2^{-\lambda}$, that a random Chung graph has expansion $\beta(\alpha)$ for every $\alpha \in [\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}]$. We show a concrete example of parameter setting in Appendix C.

A.1 Concrete Parameters of Expansion

We will be using $d = 8$. Using a numerical calculation for β_{optmial} (defined in the previous section), and the fact that we can set $\varepsilon_{\mathsf{chung}}$ small enough to ensure that β is a close approximation of β_{optmial}, we can see that expansion is quite fast (see Figure 2).

Numerical experiments show that expansion is by a factor of at least 3 for α satisfying $\delta \leq \alpha \leq 0.14$ and that $\beta_\delta^3(0.14) > 1 - 3 \cdot 0.14$ assuming $\delta \leq 0.14$; by Claim 13, expansion for weight $1 - 3 \cdot 0.14$ and above rapidly gets the weight close to 1. These statements are summarized in Fact 2 in Section 8.

B Facts about Concave Functions

Recall the definition of a concave function: F is concave if for all a, b, the graph of F on the segment $[a, b]$ does not dip below the line connecting the points $(a, F(a))$ and $(b, F(b))$. Algebraically, for all $0 \leq \lambda \leq 1$, $F(\lambda a + (1-\lambda)b) \geq \lambda F(a) + (1-\lambda)F(b)$. F is strictly concave is the inequality is strict for $0 < \lambda < 1$. Recall also that F is monotonically nondecreasing (respectively, increasing, nonincreasing, decreasing) if for all $a > b$, $F(a) \geq F(b)$ (respectively, $F(a) > F(b)$, $F(a) \leq F(b)$, $F(a) < F(b)$).

The first three claims are below standard; the last one only slightly less so.

Claim 15. *The minimum of a concave function on a line segment $[a, b]$ is reached at either a or b, and nowhere else if concavity is strict.*

Proof. Let $c \in [a, b]$ and $\lambda = (b-c)/(b-a)$. Let $m = \min(F(a), F(b))$. Then $F(c) = F(\lambda a + (1-\lambda)b) \geq \lambda F(a) + (1-\lambda)F(b) \geq \lambda m + (1-\lambda)m = m$. If the concavity is strict, then whenever $a < c < b$, $0 < \lambda < 1$ and so the first inequality is strict. □

Claim 16. *The sum of two concave functions is concave and, moreover, is strictly concave if one of the two functions is strictly concave.*

Proof. Assume F and G are concave and $H = F + G$. $H(\lambda a + (1-\lambda)b) = F(\lambda a + (1-\lambda)b) + G(\lambda a + (1-\lambda)b) \geq \lambda F(a) + (1-\lambda)F(b) + \lambda G(a) + (1-\lambda)G(b) = \lambda H(a) + (1-\lambda)H(b)$. The inequality will be strict if the inequality for F or G is strict. □

Claim 17. *Let F and G be concave nondecreasing functions. Then $F \circ G$ is a concave nondecreasing function wherever it is defined (which may not on the entire domain of G, because we do not require F to be defined on the entire range of G). (Note that for concavity of $F \circ G$, it suffices for F to be nondecreasing, and it doesn't matter whether G is nondecreasing.)*

Proof. Let $a \geq b$ be in the domain of $F \circ G$. Since G is nondecreasing, $G(a) \geq G(b)$, and thus, since F is nondecreasing, $F(G(a)) \geq F(G(b))$. Thus, $F \circ G$ is nondecreasing.

Because G is concave, $G(\lambda a + (1-\lambda)b) \geq \lambda G(a) + (1-\lambda)G(b)$. Because F is nondecreasing and concave, $F(G(\lambda a + (1-\lambda)b) \geq F(\lambda G(a) + (1-\lambda)G(b)) \geq \lambda F(G(a)) + (1-\lambda)F(G(b))$. □

Claim 18. *Let F be a concave function. Suppose $x_1 \leq x_2$ and $z \geq 0$. Let $\delta_1 = F(x_1) - F(x_1 - z)$ and $\delta_2 = F(x_2) - F(x_2 - z)$. Then $\delta_1 \geq \delta_2$.*

Proof. The intuition is simple: because F is concave, $F(x_1)$ and $F(x_2 - z)$ are both above the straight line that connects $(x_1 - z, F(x_1 - z))$ with $(x_2, F(x_2))$. If we lowered $F(x_1)$ and $F(x_2 - z)$ to this line line, we would decrease δ_1 and increase δ_2, and we would make them equal. So $\delta_1 > \delta_2$.

Algebraically, let $a = x_1 - z$, $b = x_2$, $\lambda = \frac{x_2 - x_1}{x_2 - x_1 + z}$, $\mu = \frac{z}{x_2 - x_1 + z}$. Note that $\lambda a + (1-\lambda)b = x_1$ and $\mu a + (1-\mu)b = x_2 - z$, and that $\lambda + \mu = 1$.

The concavity of F gives two inequalities:

$$F(x_1) = F(\lambda a + (1-\lambda)b) \geq \lambda F(a) + 1 - \lambda F(b) = \lambda F(x_1 - z) + (1-\lambda)F(x_2)$$

$$F(x_2 - z) = F(\mu a + (1-\mu)b) \geq \mu F(a) + 1 - \mu F(b) = \mu F(x_1 - z) + (1-\mu)F(x_2)$$

Adding them together, we get

$$F(x_1) + F(x_2 - z) \geq F(x_1 - z) + F(x_2)$$

and the result follows by subtracting $F(x_1 - z) + F(x_2 - z)$ from both sides of the inequality. □

We restate and prove Fact 1 from Section 4.1.

Fact 1. *The function gain is strictly concave on the interval* $[0, 1]$*, with* $gain(0) = gain(1) = 0$*. There is a value* $0 < \alpha_g < 1$ *that maximizes gain. The function gain (and therefore also* $gain_\delta$*) is monotonically increasing on inputs from 0 to* α_g *and monotonically decreasing on inputs from* α_g *to 1.*

Proof. *gain* is a continuous strictly concave function as a sum of two continuous concave functions β (per Condition 1) and $-\alpha$ (with β strictly concave), per claim Claim 16. It is bounded because β is bounded by 1, and a bounded continuous function reaches its maximum on the compact set [0,1]; this maximum is nonzero (because $\beta(\alpha) > \alpha$ on $(0, 1)$) and therefore not reached at 0 or 1, where *gain* is 0, so $0 < \alpha_g < 1$. It is easy to show that a violation of the monotonicity conditions on either side of α_g would imply a violation of concavity of *gain*: if $gain(y) \leq gain(x)$ for some $x < y < \alpha_g$, then $(y, gain(y))$ lies below the line connecting $(x, gain(x))$ with $(\alpha_g, gain(\alpha_g))$, as that line slopes up, since $gain(x) < gain(\alpha_g)$. Same proof works, mutatis mutandis, for the other side. □

C Concrete Results for the Filecoin Instantiation

The SDR construction of Fisch is deployed on the FIlecoin blockchain [Pro23]. It has been analyzed for a space gap of $\varepsilon_{\mathsf{space}} = 0.2$ under a depth-robustness conjecture, and has been shown to have $r_{\mathsf{hardness}} = 1 - \varepsilon_{\mathsf{hardness}} \geq \frac{1}{55}$ [GN23]. As a practical application of our results, we show that, under the same conjecture, $r_{\mathsf{hardness}} > \frac{1}{5}$ without any modification to the deployed construction — an 11-fold improvement. Moreover, adding just a few layers to G can provide further significant improvements.

In more detail, the Filecoin blockchain uses the SDR construction with a degree-8 Chung expander, connecting horizontal layers of size $n = 2^{30}$ each. It assumes (with some evidence [FBGB18] but not a proof) that its horizontal degree-6 graph is depth robust, so that a path containing 20% of the nodes exists in any given layer even if 20% of the nodes in that layer are removed (as we already discussed, we can relax this assumption, as we do not need a path, but rather a single-sink induced subgraph). The number of levels is $\ell = 11$; they all have horizontal edges, even though, as we show below, only the lowest $\ell_{\mathsf{pr}} = 8$ need to. Filecoin initialization ensures that $\delta = 0.0378$. The spacegap $\varepsilon_{\mathsf{space}}$ of interest is 0.2. We use these parameters to state the following theorem.

Theorem 5. *Suppose SPR is instantiated with a predecessor-robust graph with $\pi = 0.8$ and $\alpha_\pi = 0.2$, a Chung expander of degree 8, and parameters $\ell_{\mathsf{pr}} = 8$ and $\ell \geq 11$. Assume $\varepsilon_{\mathsf{space}} = 0.2$ and $\delta = 0.0378$. Then it has hardness ratio*

$$r_{\mathsf{hardness}} \geq \frac{2.24 + 0.93(\ell - 11)}{\ell}$$

and single-query catching probability 10%*.*

The result of this theorem is slightly better than what follows directly from the proof of Theorem 1; in particular, crucially, this theorem gives a nontrivial result for $\ell = 11$, which is the deployed instantiation, while the constants from Theorem 1 would have nothing to say until $\ell = 15$ and thus would say nothing about the deployed instantiation. Thus, theorem requires additional work, which we show in the rest of this section. Prior to this work, the best hardness ratio known for SDR with these parameters was $0.2/\ell$ [Fis19, GN23] (importantly, that hardness ratio is proven for parallel hardness, which corresponds to latency, while our result is only for sequential hardness, which corresponds to cost).

We will set $\varepsilon_{\mathsf{chung}} = 2^{-22}$, so that by Claim 14 that a random degree-8 Chung graph satisfies the expansion condition with probability at least $1 - 2^{-249}$. For the rest of this section we assume the specific choice of the Chung graph satisfies it.

As before, we set $\zeta = 1 - \varepsilon_{\mathsf{space}}/2 = 0.9$ and thus $\zeta_\delta = \zeta - \delta = 0.8622$.

In this section we use numerical estimates for values of β on specific inputs, obtained via a simple implementation of formula (4) from Section A, which finds upper and lower bounds on $\beta(\alpha)$ using binary search. With these parameter settings, we can compute

- $0.0508 < \overline{\pi} < 0.0509$
- $0.9491 < \beta(\pi) < 0.9492$
- $0.1113 < g_\pi < 0.1114$
- $0.2925 < g_{\alpha_\pi} < 0.2926$
- $0.0097 < \alpha_\delta^{\mathsf{min}} < 0.0098$
- $0.9524 < \alpha_\delta^{\mathsf{max}} < 0.9525$
- $0.3200 < \alpha_g < 0.3202$ and
- $0.3599 < \mathit{gain}(\alpha_g) < 0.3600$

In the rest of this section we prove Theorem 5.

We could simply plug in the above numbers into the proof of Theorem 1 to get results for this insantiation. We could pick $\sigma = 0.1$ and $r = 0.92$ and we would get $m_1 = 7$ from Theorem 2 (see Corollary 3), $m_2 = 3 = \beta\mathrm{count}_{0.2}(0.9)$ from Theorem 4 (thus $\ell_{\mathsf{pr}} = m_1 + m_2 = 10$), and $m_3 = 4 = \beta\mathrm{count}_{0.1}(0.92) = 4$. This would give us a total footprint of $0.92 \cdot (\ell - 14)$. What we get instead is better by about 5 levels ($0.93 \cdot 3 + 2.24 > 0.92 \cdot 5$). While the difference may seem minor, it is crucial for small ℓ and, in particular, for the deployed version of SDR, where $\ell = 11$.

In this section, we will use α_i to denote $f_i(\zeta_\delta - \rho_\ell, \ell)$ (i.e., the lower bound on the footprint of a set S of weight ζ on the bottom level ℓ). We will continue use f_i to denote the bound on the footprint of an unpebbled set T of weight α_π on level b.

We find room for improvement for these specific parameters for the following reasons:

- The proof of Theorem 1 separately counts, and skips, the maximum number of infertile levels (Theorem 2) and the maximum number of fertile levels that are not going to grow (Theorem 4). But keeping fertile levels from growing takes a lot of pebbles, which reduces not only the footprint of T, but also the footprint α_i of S. A lower α_i results in a bigger gain in the footprint of S, which increases the number of pebbles necessary to make an infertile level, and therefore reduces the number of infertile levels. In other words, the existing proof does not take advantage of the fact that more levels for Theorem 4 means fewer levels for Theorem 2 and vice versa.
- The proof of Theorem 4 ignores fertile levels whose footprints dip below α_π before rebounding and growing.
- The proof of Theorem 1 ignores the footprint of a growing fertile level until it reaches weight r.

C.1 Number of Infertile Levels as a Function of Pebble Arrangements

Most of the work in this section is simply in applying the general results in Section 5 to the specific parameters of Theorem 5. However, Claims 22, 23, and 24 are new, and address the relationship between footprints, pebbles spent, and the number of fertile levels.

Claim 19. *For all i, $\alpha_i \geq 0.0622 > \overline{\pi}$ and therefore for every infertile level m, $gain_\delta(\alpha_m) \geq g_\pi > 0.1113$.*

Proof. By Claim 5, α_i never falls below $\zeta_\delta - \rho = 0.0622$ and $0.0622 > \overline{\pi}$, because $gain_\delta(0.0622) > gain_\delta(0.8)$. □

Thus, the simple case of Theorem 2—namely, the one given by Lemma 1—applies and we have the following corollary to Theorem 2. Note that we get at most 7 for the number of infertile levels (we argue that this is tight in Appendix C.4), while the best previously known bound was 10 [Fis19, GN23].

Corollary 3. *For the parameter settings in Theorem 5, the number of infertile levels is at most 7 and the following holds for any level m:*

if number of infertile levels below level m is	*then maximum weight of black pebbles $\rho_{1\ldots m}$ at level m and above is at most*
1	*0.7378*
2	*0.6265*
3	*0.5152*
4	*0.4039*
5	*0.2926*
6	*0.1813*
7	*0.0700*

It is helpful to have the following variant of Lemma 1.

Claim 20. *Let m be the highest infertile level (assuming one exists).*

$$\sum_{i=m}^{\ell} gain_\delta(\alpha_i) < 0.8492$$

Proof. By Claim 5, $\sum_{i>m} gain_\delta(\alpha_i) \le \alpha_m - \zeta + \rho + \delta$. Add $gain_\delta(\alpha_m)$ to both sides of the inequality, and recall that $\alpha_m + gain_\delta(\alpha_m) = \beta(f_m) - \delta < \beta(\pi) - \delta$ because $\overline{\pi} < f_m < \pi$ because $\alpha_i > \overline{\pi}$ for all i by Claim 19 and level m is fertile so $\alpha_i < \pi$. □

Claim 21. *If i is the lowest fertile level, then $gain_\delta(\alpha_i) \ge 0.0313$.*

Proof. If $i = \ell$, then $\alpha_i \le \zeta_\delta = 0.8622$. Else, the level below i is infertile, and thus $\alpha_{i+1} < \pi$, so $\alpha_i < \beta_\delta(\pi) < 0.9114$. Because $gain_\delta$ monotonically decreases above π, and $\alpha_i \ge \pi$ because i is fertile, we have $gain_\delta(\alpha_i) > gain_\delta(0.9114) > 0.0313$. □

The following claim shows that if the number of infertile levels is maximum possible, then fertile levels cannot be extinguished, because $gain_\delta(\alpha_\pi) > 0.2107$. This shows that the maximum number of levels for Theorem 2 leaves no levels for Theorem 4.

Claim 22. *Let b be the lowest fertile level. If there are 7 infertile levels, then for every level $m < b$ above b, $\rho_m < 0.2107$.*

Proof. First, note that $gain_\delta(\alpha_m) < 0.1501$. Indeed, this is automatically true for fertile m, because for a fertile m, $gain_\delta(\alpha_m) \le g_\pi < 0.1114$. If this is false for some infertile m, then, since there are 7 infertile levels total, taking m' to be the highest infertile level, we have $\sum_{i=m'}^{\ell} gain_\delta(\alpha_i) \ge gain_\delta(\alpha_b) + 0.1501 + 6 \cdot g_\pi > 0.0313 + 0.1501 + 6 \cdot 0.1113 = 0.8492$ (by Claim 19 and 21), which contradicts Claim 20.

There are two regions of [0,1] where $gain_\delta(\alpha) \le 0.1501$: one requires that $0 \le \alpha < 0.0703$ and the other requires that $0.7418 < \alpha \le 1$. By Claim 5,

$$\alpha_m \ge \alpha_b + gain_\delta(\alpha_b) - \rho = \beta_\delta(\alpha_b) - 0.8 \ge \beta_\delta(\pi) - 0.8 > 0.0703\,,$$

so we must have $\alpha_m > 0.7418$.

Note that $\beta_\delta(\alpha_{m+1}) < \alpha_\delta^{\mathsf{max}}$ (else $gain_\delta(\beta_\delta(\alpha_{m+1})) \le 0$, which contradicts Claim 5). Since $\alpha_m = \beta_\delta(\alpha_{m+1}) - \rho_m$, we have $\rho_m < \alpha_\delta^{\mathsf{max}} - 0.7418 < 0.9525 - 0.7418 < 0.2107$. □

The next claim shows how a small footprint α_m (which can happen when a lot of pebbles are used to extinguish a fertile level) reduces the number of infertile levels.

Claim 23. *If there is m with $\alpha_m \le 0.5015$, then there are at most 5 infertile levels.*

Proof. Suppose there are 6 or more infertile levels. If at least four of those are below m, then by Claim 5 and Claim 19

$$\alpha_m \ge \zeta - \delta - \rho + 4 \cdot g_\pi = 0.0622 + 0.1113 \cdot 4 > 0.5015\,,$$

which is a contradiction. Thus, at most three infertile levels are below m, so there are at least two levels above m.

The main idea of the proof is to show that the gains of levels m and $m-1$ are too high. We will consider two cases: $\alpha_m > 0.2023$ and $\alpha_m \le 0.2023$.

Suppose $\alpha_m > 0.2023$. By Claim 15, because *gain* is concave (Fact 1), we know $gain_\delta(\alpha_m) \le \min(gain_\delta(0.2023), gain_\delta(0.5015)) > 0.2804$ To have $\alpha_m \le 0.5015$, we had to place black pebbles of weight at least $\zeta_\delta - 0.5015 = 0.8622 - 0.5015 = 0.3607$ at level m or below (by Claim 5), which means that the weight of the black pebbles above level m is at most $\rho - 0.3607 = 0.4393$. Then we know

$$\alpha_{m-1} > \beta_\delta(\alpha_m) - 0.4393 > \beta_\delta(0.2023) - 0.4393 > 0.0567$$

and

$$\alpha_{m-1} \le \beta_\delta(0.5015) < 0.7820\,,$$

so $gain_\delta(\alpha_{m-1}) \ge \min(gain_\delta(0.0567), gain_\delta(0.7820)) > 0.1236$. Note also that level $m-1$ is infertile, because $\alpha_{m-1} < 0.7820 < 0.8 = \pi$.

Taking m' to be the highest infertile level, we have by Claim 19

$$\begin{aligned}\sum_{i=m'}^{\ell} gain_\delta(\alpha_i) &\ge 4 \cdot g_\pi + gain_\delta(\alpha_m) + gain_\delta(\alpha_{m-1}) \\ &> 0.4452 + 0.2804 + 0.1236 = 0.8492\,,\end{aligned}$$

which contradicts Claim 20. This concludes the first case.

Now consider the second case: suppose $\alpha_m \le 0.2023$. By Claims 5 and 19, because there are at least 5 infertile levels below the highest fertile level m', $\alpha_{m'} \ge 0.0622 + 5 \cdot 0.1113 = 0.6187 > 0.5015$. We thus know that there exists at least one level above m for which $\alpha_i > 0.5015$. Let $i < m$ be the lowest such level. Then $\alpha_{i+1} > 0.2023$ (because $\beta_\delta(0.2023) < 0.5015$) but $\alpha_{i+1} \le 0.5015$ by the definition of i, and thus we can apply the previous case to $m = i + 1$. □

Finally, we show that extinguishing even one fertile level (which costs $g_{\alpha_\pi} > 0.2925$ pebble weight) while keeping at least six infertile levels reduces the number of available black pebbles.

Claim 24. *Suppose level $b \le \ell - 6$ is infertile, and at least five levels below it are also infertile. If there is a level $m \ge b$ with $\rho_m > 0.2925$, then the weight $\rho_{1\ldots b-1}$ of black pebbles above b is at most* 0.0607*.*

Proof. First consider the case $m > b$.

Note that level m is infertile because $\alpha_m < 1 - 0.2925 < 0.8 = \pi$. Because level b is infertile, $0.8 > \alpha_b$, and thus by Claims 5, 19, and 21 (because at least four levels below b, besides level m, are infertile; at there is at least one more level that is either fertile or infertile)

$$\begin{aligned} 0.8 > \alpha_b &= \zeta_\delta + \sum_{i>b} gain_\delta(\alpha_i) - \rho_{\ell\ldots b} \\ &\geq 0.8622 + 0.0313 + 0.1113 \cdot 4 + gain_\delta(\alpha_m) - \rho_{\ell\ldots b}\,. \end{aligned}$$

Level m has $\alpha_m = \beta_\delta(\alpha_{m+1}) - 0.2925$. We know $\beta_\delta(\alpha_{m+1}) < \alpha_\delta^{\mathsf{max}} < 0.9525$ (because otherwise $gain_\delta(\beta_\delta(\alpha_{m+1}) \leq 0$, which contradicts Claim 5), so we have $\alpha_m < 0.66$. Since there are at least six infertile levels, by Claim 23, $\alpha_m > 0.5015$. Because *gain* is monotonically decreasing on inputs in the range from 0.5015 to 0.66 by Fact 1, $gain_\delta(\alpha_m) > gain_\delta(0.66) > 0.2006$. Thus, we have $\rho_{\ell\ldots b} > 0.7393$ and $\rho_{1\ldots b-1} < \rho - 0.7393 = 0.0607$.

If $m = b$, let $m' > b$ be the highest infertile level below m. It has at least four infertile levels below it, so by Claims 5 and 21

$$\begin{aligned} 0.8 > \alpha_{m'} &= \zeta_\delta + \sum_{i>m'} gain_\delta(\alpha_i) - \rho_{\ell\ldots m'} \\ &\geq 0.8622 + 0.1113 \cdot 4 - \rho_{\ell\ldots m'}\,, \end{aligned}$$

and so $\rho_{\ell\ldots m'} \geq 0.5074$. Thus, $\rho_{\ell\ldots b} \geq 0.5074 + \rho_b > 0.7999$, so $\rho_{1\ldots b-1} < \rho - 0.7999 = 0.0001 < 0.0607$. □

C.2 Footprints For Specific Pebble Arrangements

Start with an unpebbled set T of weight $f_b = \alpha_\pi = 0.2$ on level b. In this section, we show lower bounds on the total footprint of b in several different situations. These situations do not cover all possibilities, but they turn out to be sufficient for the final proof in Appendix C.3. The specific situations addressed in this section are:

- When $b \geq 4$ and $\rho_{b-1\ldots1} \leq 0.07$ (Claim 25)
- When $b \geq 5$ and $\rho_{b-1\ldots1} \leq 0.30$ (Claim 26)
- When $b \geq 6$ and $\rho_{b-1\ldots1} \leq 0.44$ (Claim 27)
- When $b \geq 8$ and T is viable for at least three levels (Claim 28)
- When $b \geq 8$ and $\rho_{b-1} + \rho_{b-2} \leq 0.36$ and $\rho_{b-1\ldots1} \leq 0.8$ (Claim 29)
- When $b \geq 8$ and $\rho_{b-1} \leq 0.1525$, $\rho_{b-1} + \rho_{b-2} \leq 0.73$, and $\rho_{b-1\ldots1} \leq 0.8$ (Claim 30)

The first four of these claims simply apply the results of Sections 6 and 7 to the specific parameters of Theorem 5. The last two are new, because they deal footprints of fertile sets that may lose viability for a few levels and then regain it. These claims require calculations of the functions β, *gain*, and ϕ. We do not show these calculations explicitly—they are done by straightforward code that computes the function β for the Chung expander.

Claim 25. *Suppose $b \geq 4$, $f_b = 0.2$, and the total weight $\rho_{b-1\ldots1}$ of black pebbles above level b is at most 0.07. The total footprint is at least $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) > 2.24 + 0.93 \cdot (b-4)$.*

Proof. Applying Theorem 3, we know $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) \geq \phi_{f_b}(0.07, 0, \ldots, 0) > 2.24 + 0.93 \cdot (b-4)$. □

Claim 26. *Suppose $b \geq 5$, $f_b = 0.2$, and the total weight $\rho_{b-1\ldots1}$ of black pebbles above level b is at most 0.3. Then the total footprint is at least $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) > 2.54 + 0.94 \cdot (b-5)$.*

Proof. Applying Theorem 3, we know $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) \geq \phi_{f_b}(0.3, 0, \ldots, 0) > 2.54 + 0.94 \cdot (b-5)$. □

Claim 27. *Suppose $b \geq 6$, $f_b = 0.2$, and the total weight $\rho_{b-1\ldots1}$ of black pebbles above level b is at most 0.44. The total footprint is at least $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) > 2.49 + 0.93 \cdot (b-6)$.*

Proof. Applying Theorem 3, we know $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) \geq \phi_{f_b}(0.44, 0, \ldots 0) > 2.49 + 0.93 \cdot (b-6)$. □

Claim 28. *Suppose T is an unpebbled subset of level $b \geq 8$ of weight $f_b = \alpha_\pi = 0.2$ that is viable for at least 3 levels. Then the total footprint of T is at least $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) > 3.4 + 0.94 \cdot (b-8)$.*

Proof. By Corollary 2, $\rho_{b\ldots b-2} + \beta_\delta(f_{b-2}) \geq \min(0.2 + 3 \cdot gain_\delta(0.2), \beta_\delta^3(0.2)) = \beta_\delta^3(0.2) > 0.9037$. Therefore, $\beta_\delta(f_{b-2}) - \rho_{b-3\ldots1} = \beta_\delta(f_{b-2}) + \rho_{\ell\ldots b-2} - \rho > 0.1037$. Thus, by Theorem 3. the footprint of T is at least $0.2 \cdot 2 + \phi_{0.2}(\rho_{b-3\ldots1}, \underbrace{0, \ldots, 0}_{b-4}) = 0.6 + \phi_{0.1037}(\underbrace{0, \ldots 0}_{b-4}) \geq 0.6 + 2.8 + 0.94 \cdot (b-8)$. □

Claim 29. *Suppose $b \geq 8$, $f_b = 0.2$, $\rho_{b-1} + \rho_{b-2} \leq 0.36$, and the total weight $\rho_{b-1\ldots1}$ of black pebbles above level b is at most 0.8. Then the total footprint is at least $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) > 3.07 + 0.93 \cdot (b-8)$.*

Proof. We have $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) \geq \phi_{f_b}(\rho_{b-1}, \rho_{b-2}, \rho_{b-3\ldots1}, 0, \ldots, 0) \geq \phi_{f_b}(\rho_{b-1}, \rho_{b-2}, 0.8 - \rho_{b-1} - \rho_{b-2}, 0, \ldots, 0)$ (by applying Lemma 3 $b-4$ times followed by Claim 2)

For ease of notation, fix $b = 8$ for now. We will provide a lower bound for $\phi_{f_b}(\rho_7, \rho_6, 0.8 - \rho_7 - \rho_6, 0, 0, 0, 0)$, where $\rho_7 + \rho_6 \leq 0.36$. The $f_1, \ldots, f_8$ values discussed in the rest of the proof are with respect to this calculation of ϕ.

Since $\beta_\delta(0.2) - 0.36 \leq f_7 \leq \beta_\delta(0.2)$, and *gain* is concave, we have

$$\begin{aligned} gain_\delta(f_7) &\geq \min(gain_\delta(\beta_\delta(0.2)), gain_\delta(\beta_\delta(0.2) - 0.36)) \\ &= gain_\delta(\beta_\delta(0.2) - 0.36) > 0.2389\,. \end{aligned}$$

Using Claim 3, $f_6 = \beta_\delta(0.2) + gain_\delta(f_7) - (\rho_6 + \rho_7)$, and thus

$$\beta_\delta(0.2) + gain_\delta(\beta_\delta(0.2) - 0.36) - 0.36 < f_6 \leq \beta_\delta(f_7) \leq \beta_\delta(\beta_\delta(0.2))\,.$$

Therefore, $0.3715 < f_6 < 0.7766$.

Because *gain* is concave, $gain_\delta(f_6) \geq \min(gain_\delta(0.3715), gain_\delta(0.7766)) = gain_\delta(0.7766) > 0.1271$. Thus, using Claim (3), $f_5 = \beta_\delta(0.2) + gain_\delta(f_7) + gain_\delta(f_6) - 0.8 \geq \beta_\delta(0.2) + gain_\delta(\beta_\delta(0.2) - 0.36) + gain_\delta(0.7766) - 0.8 > 0.0586$.

Thus, by Claim 2, $f_5+f_4+3+f_2+f_1 = \phi_{f_5}(0,0,0,0) > \phi_{0.0586}(0,0,0,0) > 2.37$. And by Lemma 3 and Claim 2, $f_8 + f_7 + f_6 = \phi_{0.2}(f_7, f_6) \geq \phi_{0.2}(f_7 + f_6, 0) \geq \phi_{0.2}(0.36, 0) > 0.7$, giving us a total of $2.37 + 0.7 = 3.07$.

If $b > 8$, we simply replace $\phi_{0.0586}(0,0,0,0)$ with $\phi_{0.0586}(\underbrace{0,\ldots 0}_{b-4}) > 2.37 + 0.93 \cdot (b - 8)$. □

Claim 30. *Suppose $b \geq 8$, $f_b = 0.2$, $\rho_{b-1} \leq 0.1525$, $\rho_{b-1} + \rho_{b-2} \leq 0.73$, and the total weight $\rho_{b-1\ldots1}$ of black pebbles above level b is at most 0.8. Then the total footprint is at least $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) > 3.17 + 0.94 \cdot (b - 8)$.*

Proof. We have $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) \geq \phi_{f_b}(\rho_{b-1}, \rho_{b-2}, \rho_{b-3\ldots1}, 0, \ldots, 0) \geq \phi_{f_b}(\rho_{b-1}, \rho_{b-2}, 0.8 - \rho_{b-1} - \rho_{b-2}, 0, \ldots, 0)$ (by Applying Lemma 3 $b-4$ times followed by Claim 2).

For now, for ease of notation, we will fix $b = 8$ and provide a lower bound for $\phi_{f_b}(\rho_7, \rho_6, 0.8 - \rho_7 - \rho_6, 0, 0, 0, 0)$, where $\rho_7 \leq 0.1525$ and $\rho_7 + \rho_6 \leq 0.73$. The $f_1 \ldots f_8$ values discussed in the rest of the proof are with respect to this calculation of ϕ.

Since $\beta_\delta(0.2) - 0.1525 \leq f_7 \leq \beta_\delta(0.2)$, we have $0.34 < f_7 < 0.4926$. Since $gain_\delta$ is decreasing above $\beta_\delta(0.2) - 0.1525 > \alpha_g$, we have

$$gain_\delta(f_7) \geq gain_\delta(\beta_\delta(0.2)) > 0.2839\,.$$

Using Claim 3, $f_6 = \beta_\delta(0.2) + gain_\delta(f_7) - (\rho_6 + \rho_7)$, and thus

$$\beta_\delta(0.2) + gain_\delta(\beta_\delta(0.2)) - 0.73 < f_6 \leq \beta_\delta(f_7) \leq \beta_\delta(\beta_\delta(0.2))\,.$$

Therefore, $0.0465 < f_6 < 0.7766$.

Because $gain_\delta$ is concave, $gain_\delta(f_{b-2}) \geq \min(gain_\delta(0.0465), gain_\delta(0.7766)) = gain_\delta(0.0465) > 0.1016$. Thus, using Claim 3, $f_5 = \beta_\delta(0.2)+gain_\delta(f_7)+gain_\delta(f_6)-0.8 \geq \beta_\delta(0.2) + gain_\delta(\beta_\delta(0.2)) + gain_\delta(0.0464) - 0.8 > 0.0782$.

Thus, by Claim 2, $f_5 + f_4 + 3 + f_2 + f_1 = \phi_{f_5}(0,0,0,0) > \phi_{0.0782}(0,0,0,0) > 2.59$. In addition, $f_8 + f_7 + f_6 > 0.2 + 0.3399 + 0.0464 > 0.58$, for a total of at least 3.17.

If $b > 8$, we simply replace $\phi_{0.0782}(0,0,0,0)$ with $\phi_{0.0782}(\underbrace{0,\ldots 0}_{b-4}) > 2.59 + 0.94 \cdot (b - 8)$. □

C.3 Putting the Proof of Theorem 5 Together

We now prove Step 3.5 described in Section 3.2, which suffices for proving Theorem 5.

Lemma 5. *There is a fertile level $b \geq \ell - 7$ with the following property. Let T be an unpebbled subset of this level with weight at least 0.2. The total footprint of T is at least $2.24 + 0.93 \cdot (\ell - 11)$.*

Proof. Let m_1 be the lowest fertile level. We know $m_1 \geq \ell - 7$, because there are at most 7 infertile levels by Corollary 3.

If $m_1 = \ell - 7$, then there are at least 7 infertile levels below m_1, and thus the weight of black pebbles above m_1 is at most 0.07 by Corollary 3, and thus we can set $b = m_1$ and apply Claim 25 to bound the total footprint.

If $m_1 = \ell - 6$, then there are at least 6 infertile levels below m_1, and thus the weight of black pebbles above m_1 is at most 0.1813 by Corollary 3, and thus we can set $b = m_1$ and apply Claim 26 to bound the total footprint.

If $m_1 = \ell - 5$ or $m_1 = \ell - 4$, then there are at least 4 infertile levels below m_1, and thus the weight of black pebbles above m_1 is at most 0.4039 by Corollary 3, and thus we can set $b = m_1$ and apply Claim 27 to bound the total footprint.

If $m_1 \geq \ell - 3$, the proof gets harder, because now the adversary may have enough pebbles above m_1 to stop the growth of T on m_1 completely (since $\beta_\delta(\alpha_\pi) \approx 0.4925$, pebble weight 0.4925 right above m_1 suffices, and Corollary 3 cannot rule it out). We will have to proceed by cases: in some cases, there won't be enough pebbles immediately above m_1, and T will grow, and in other cases, there will be many pebbles immediately above m_1, but then second or third lowest fertile level will grow, because there won't be enough pebbles to stop the growth above those levels.

If $m_1 \geq \ell - 3$ and there are exactly 7 infertile levels (there cannot be more by Corollary 3), then for each $i < \ell - 3$, each $\rho_i \leq 0.2107 < g_\pi$ by Claim 22, and thus, by simple induction, m_1 can never be extinguished, so we can apply Claim 28 to bound the total footprint.

If $m_1 \geq \ell - 3$ and there are 6 or fewer infertile levels, let m_2 be the second lowest fertile level. Note that $m_2 \geq \ell - 7$ because there are at most 6 infertile levels.

- If $m_2 = \ell - 7$. We will do a proof by cases, depending on how concentrated the pebbles are above m_1. If no level i such that $m_2 < i < m_1$ has $\rho_i > g_\pi > 0.2925$, set $b = m_1$. By simple induction, m_1 is viable for at least 3 levels, so we can apply Claim 28 to bound the total footprint. Else, we know that the weight of black pebbles above level m_2 is at most 0.0607 by Claim 24 (where b in Claim 24 is set to $m_2 + 1 = \ell - 6$; note that there are six levels below $\ell - 6$ and all but m_1 are infertile). So we set $b = m_2$ and can apply Claim 25 to bound the total footprint.

- If $m_2 = \ell - 6$, then there are at least 5 infertile levels below m_2, and thus the weight of black pebbles above m_2 is at most 0.2926 by Corollary 3, and thus we can set $b = m_2$ and apply Claim 26 to bound the total footprint.

- If $m_2 = \ell - 5$, then there are at least 4 fertile levels below m_2, and thus the weight of black pebbles above m_2 is at most 0.4039 by Corollary 3, and thus we can set $b = m_2$ and apply Claim 27 to bound the total footprint.

- If $m_2 \geq \ell - 4$, we again have the problem that the adversary has enough pebbles to stop the growth of T from m_2. We consider two cases, with two subcases each.

  - $m_1 \geq m_2 + 2$. We will show that either m_1 or m_2 will grow, since there is not enough pebble weight to stop the growth of both. The cases will focus on how much pebble weight there is between m_1 and m_2. Specifically, if $\rho_{m_2 \ldots m_1} \geq 0.36$, set $b = m_2$. We know the weight of black pebbles above m_2 is at most $\rho - \rho_{m_2 \ldots m_1} \leq 0.44$, and since $m_2 > \ell - 5$, we can apply Claim 27

to bound the total footprint. And if $\rho_{m_2 \ldots m_1} < 0.36$, then in particular $\rho_{m_1-2 \ldots m_1-1} < 0.36$, so we set $b = m_1$ and apply Claim 29 to bound the total footprint.

- $m_1 = m_2 + 1$. Since we have two fertile levels in a row, m_1 has a chance to grow for at least one level. Specifically, we know $\rho_{m_2} \le 0.1525$ because $\rho_{m_2} = \beta(m_1) - m_2 \le \alpha_\delta^{\mathsf{max}} - \pi < 0.1525$ (since $\beta(m_1) < \alpha_\delta^{\mathsf{max}}$ by Claim 4). This growth can still be stopped, but it will require a lot of pebbles in the level above m_2. Specifically, if $\rho_{m_2} + \rho_{m_2-1} \le 0.73$, then set $b = m_1$ and apply Claim 30 to bound the total footprint. Else, there are at most five infertile levels by Claim 23, because $\alpha_{m_2-1} \le 1 - \rho_{m_2-1} - \delta \le 1 - (0.73 - 0.1525) - \delta = 0.3847$. Thus, there is a fertile level m_3 such that $\ell - 7 \le m_3 \le m_2 - 1$, and the weight of black pebbles above m_3 is less than $0.8 - \rho_{m_2-1} + \rho_{m_2} < 0.8 - 0.73 = 0.07$. Set $b = m_3$. Claim 25 applies to bound the total footprint.

This concludes the proof of Lemma 5. □

C.4 On Optimality of the Result

Suppose the adversary places its black pebbles as follows: $\rho_\ell = 0.0623$, $\rho_{\ell-1} = \cdots = \rho_{\ell-6} = 0.1114$, $\rho_{\ell-7} = 0$, $\rho_{\ell-8} = 0.0693$, and $\rho_{\ell-9 \ldots 1} = 0$. Then the bottom seven levels are infertile; the remaining ones are fertile. If we set $b = \ell - 7$, we get $f_{\ell-7} = 0.2$, $f_{\ell-8} \approx 0.42$, $f_{\ell-9} \approx 0.73$, and $f_{\ell-10} \approx 0.89$, for a total of about 2.24 when $\ell = 11$. Setting $b = 3, 2$, or 1 gives smaller results. If we have more levels, $f_{\ell-11} \approx 0.94$ and $f_i \approx 0.95$ for $i < \ell - 11$.

Thus, arguments that are based on the same framework of simply counting sizes (i.e., looking at vertical expansion and subtracting pebbles) for this construction are unlikey to overcome the $2.24 + 0.95 \cdot (\ell - 11)$ bound, which essentially matches the result of Theorem 5. That doesn't mean the result can't be improved—perhaps other proof frameworks than the one in Section 3 are possible. In particular, it may be possible to reason about single-node expansion, or to measure footprint growth via horizontal edges, or take into account pebble positions, or use ancestor robustness of different size sets.