Proofs of Space with Maximal Hardness

Leonid Reyzin

Boston University[⋆]
https://www.cs.bu.edu/fac/reyzin
reyzin@bu.edu

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 computation in order to satisfy the verifier.

We present the first proof space 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.

1 Introduction

In a proof of persistent space [6], 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, P takes a small instance identifier x, and generates a very large output $y = f(x)$. (Depending on the flavor of the proof of space, P may also provide to V a commitment to y and a proof that the computation of f was correct or at least close to correct.) V then periodically queries random portions of y, which P returns together with the proof of their correctness. 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 hard for P to come up with portions of y in response to the queries of V.

To make this notion more precise, we have to address what "less than all of y" and "hard" 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 then $|y|$, perhaps answering 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 "hard"; but above $(1 - \varepsilon_{\mathsf{space}}) \cdot |y|$ storage, the proof of space provides no guarantees.

[⋆] Work done 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 [3].

The construction of Pietrzak [12, Lemma 8] is the first one to provide security against an arbitrarily small space gap, but only in a theoretical sense, because the complexity of f grows too high as the space gap shrinks, and is problematic even for a space gap of $\frac{1}{2}$ [8, Section 1.1]. The SDR construction of Fisch [8] is the first one to do so for practical parameter values and is deployed in practice for a space gap of 0.2 [10].

Because x is short, answering the queries of V is never harder than simply recomputing $y = f(x)$. We will use the term "relative hardness" $r_{\mathsf{hardness}} \leq 1$ to denote the ratio between the hardness of answering the queries of V by a cheating prover who uses less than $(1 - \varepsilon_{\mathsf{space}}) \cdot |y|$ storage, and the hardness of f. We will also use the term "hardess gap" $\varepsilon_{\mathsf{hardness}} = 1 - r_{\mathsf{hardness}}$.

We have not yet addressed what "hardness" actually means. Two notions have found use in practice [10, Section 1.2]: monetary cost and latency. If the cost of recomputing portions of y in response to V's queries is higher than cost of storing $(1 - \varepsilon_{\mathsf{space}}) \cdot |y|$ bits and simply looking up the answers, then a prover incentivized by money will choose to use storage. Alternatively, if the time required to recompute portions of y is longer than V is willing to wait, then a prover who does not use storage will simply fail. While real-life cost and latency are difficult to estimate, as they depend on a multitude of implementation-specific factors, they are reasonably well approximated by the notions of sequential time and parallel time, respectively.

For both cost and and latency, the closer relative hardness r_{hardness} is to 1, the better. When hardness is measured in terms of cost, a bigger r_{hardness} lowers initializations costs for the prover and/or decreases the frequency of queries from the verifier. When hardness is measured in terms of latency, a bigger r_{hardness} reduces the time needed for initialization and/or increases V's willingness to wait, potentially allowing P to use cheaper hardware. When proofs of space are used for useful storage, as in constructions of [4,12,9,7], these gains directly translate into reduced costs for storage and retrieval.

1.1 Our Contribution

We show a practical proof of space that can achieve hardness ratio r_{hardness} arbitrarily close to 1. Moreover, we do so while also maintaining the arbitrarily small space gap $\varepsilon_{\mathsf{space}}$ shown by Fisch [8]. Our result is for sequential time (corresponding to cost); achieving the same for parallel time (corresponding to latency) remains an open problem.

In contrast to our result, the result of Fisch achieves hardness for parallel time, but r_{hardness} goes to zero as the space gap $\varepsilon_{\mathsf{space}}$ decreases. The only known proof of space that could theoretically achieve r_{hardness} arbitrarily close to 1 and $\varepsilon_{\mathsf{space}}$ arbitrarily close to 0 for parallel time is the construction of Pietrzak [12] generalized to larger and more depth-robust graphs (as presented in [12], $r_{\mathsf{hardness}} = \frac{1}{4}$), but such generalization will even further reduce the practicality of the construction.

Our construction of f, which we call SPR, is quite practical — in fact, it is a minor relaxation of the SDR construction by Fisch. Our main contribution is in the new analysis techniques. The analysis techniques we develop are more general than the techniques used in prior analyses of graph-based proof of space constructions (as they do not rely on the particular parameters of the expansion and depth-robustness) and, at the same time, provide better concrete results when instantiated with specific parameters. We emphasize that these techniques are not only of theoretical interest; we apply them to demonstrate that the relative hardness r_{hardness} for SDR as deployed by the decentralized storage blockchain Filecoin [11] is 11 times better than previously proven, and that adding just a little more complexity to the graph provides further dramatic improvements.

1.2 Technical Overview

We now elaborate on the construction and analysis.

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 (with the exception of [1]), our construction SPR uses a directed acyclic graph G with a single source to represent f, 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.

Assume that a malicious prover P^* stores labels of only some of the nodes in G.[2] Then, when 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. Sequential time corresponds to the total number of pebbles placed in order to reach v, and parallel time corresponds to the longest path on those pebbles.

A malicious prover can cheat somewhat when computing $y = f(x)$, by labeling some nodes with incorrect, easy to compute, values. This cheating cannot cover too many nodes, as the verifier spot checks the computation during the initialization phase. For P^*, incorrectly computed nodes correspond to additional pebbles on G. Pebbles of this "cheating" type are called "red" [6], in contrast to pebbles that correspond to storage, which are "black".

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

The SDR ("Stacked Depth-Robust") graph used for the computation of f in [8] is built a follows. Take ℓ layers of n nodes each; number them from 1 to ℓ

[2] As shown in [12, Sections 5, 7], for proofs of parallel time, it is possible to prove that a malicious prover who stores information other than random oracle outputs — for example, functions of those outputs — cannot do better; unfortunately, no proof of this fact is not known for sequential time.

top to bottom. Level i has edges going to level $i+1$ so as to form an expander when viewed backwards; that is, a set of nodes on level $i+1$ of size αn has $\beta(\alpha) \cdot n$ predecessors on level i, where $\beta(\alpha) > \alpha$ is a function that grows quickly, particularly for small values of α. Each level also has horizontal (left-to-right) edges that form a graph of with the following property: a path of length $\alpha_\pi \cdot n$ exists in any subgraph of size $\pi \cdot n$ (this is known as a $(1-\pi, \alpha_\pi)$-depth-robust graph). The horizontal graphs have a single source on the left; the label of the top left node is x and the labels of the entire bottom level are y.

Our construction SPR ("Stacked Predecesor-Robust") is a slight generalization of SDR: we can relax depth robustness to predecessor robustness, we do not need horizontal edges on all levels, and we are agnostic to the specific expander and depth-robustness parameters; see Section 2.

The technical heart of this paper develops techniques for lowerbounding footprint sizes in such graphs. We show the applicability of these techniques in both the asymptotic regime (where we prove that the footprint size can get arbitrarily close to $|G|$ to get r_{hardness} arbitrarily close to 1) and the concrete regime (where we provide an 11-fold improvement over the previous analysis of a deployed scheme, and show that r_{hardness} can be easily improved further).

2 Definition and Construction

2.1 The Graph SPR

Please refer to Section 1.2 for the explanation of the construction; here we only fill in the details.

Both SDR [8] and our construction SPR are graphs G of ℓ levels of n nodes each. We follow the numbering in [8]: the top level is 1, the bottom is ℓ, with edges going left-to-right in each level (for depth-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 path of length $\alpha_\pi n$. This guarantee needs to apply to all levels.

In SPR, we relax this 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 [2]: any subgraph of a given level of size $\pi \cdot n$ contains single-sink subgraph of size $\alpha_\pi n$ (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 5.1.

The levels above the lowest ℓ_{pr} need no horizontal edges, except level 1, which needs an edge from the leftmost node labeled x to every node.

The expansion guarantee is the following: any set of $\alpha \cdot n$ nodes in a given level has, collectively, $\beta(\alpha) \cdot n$ predecessors (for constant α and sufficiently large n). In contrast to SDR, which is analyzed for a specific β (from the degree-8 Chung expander), our analysis works for general β. We only require 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 emphasize that SDR is a special case of SPR, and our analysis works for SDR as well.

2.2 Initialization

As in most proof of space schemes, 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 labels, which P reveals together with the labels of their predecessors; V verifies that the decommitments are correct and that the label of the requested node is correctly computed from the 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, which can be done quickly, because the graph is almost the same layer-to-layer, so entire columns of nodes can be queried and decommitted at once.[3]

From now on, we will assume that initialization has succeeded: that is V accepted, 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.

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 [6,13]. Note that different definitions of proofs of space highlight different parameters in parentheses; we avoid the positional parenthetical notation to avoid confusion.

[3] 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.

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 ratio r_{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 $r_{\mathsf{hardness}} \cdot N$ nodes*[4] *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^* may not grow above $r_{\mathsf{hardness}} \cdot N$, as it may be shared among all the queried nodes).

3 Main Results

Theorem 1. *Fix $\varepsilon_{\mathsf{space}} > 0$. For any constant $r < 1$, there is a setting of constants δ, ℓ_{pr}, and n in the SPR construction, and a constant m, such that for any number of layers ℓ the SPR construction has hardness ratio*

$$r_{\mathsf{hardness}} \geq \frac{r(\ell - m)}{\ell}$$

(which approaches r as ℓ grows) and single-query catching probability $p_{\mathsf{hard}} \geq \varepsilon_{\mathsf{space}}/2$.

Sections 4–9 are dedicated to proving Theorem 1.

Theorem 2. *For SPR instantiated with the degree-8 Chung expander, m in Theorem 1 is at most linear in*

$$\frac{1}{\beta(\alpha_\pi) - \alpha_\pi} + \frac{1}{\beta(\pi) - \pi}$$

plus an amount that is logarithmic in the inverses of $\varepsilon_{\mathsf{space}}, \alpha_\pi, 1-\pi$ and $1-r$.

A description of Chung expanders and a proof of this theorem are in Section A.

[4] We define r_{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 r_{hardness} arbitrarily close to 1. It is also possible to use duplicate hash inputs simply to make hash computation time at each node the same (as is done in the Filecoin implementation [11]). Whether using duplicate hash inputs or increasing the number of layers in our construction results in a tighter cost ratio in practice depends on specific parameters and implementation details.

Theorem 3. *Suppose SPR is instantiated with the degree-8 Chung expander, a predecessor-robust graph with $\pi = 0.8$ and $\alpha_\pi = 0.2$, $\ell_{\mathsf{pr}} = 8$, 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 parameters in Theorem 3 are taken from the deployed instantiation on the Filecoin blockchain [11]. They are slightly better than what follows directly from the proof of Theorem 2; in particular, crucially, they give a nontrivial result for $\ell = 11$, which is the deployed instantiation, while the constants from Theorem 2 would have nothing to say until $\ell = 15$ and would say nothing about the deployed instantiation. The proof of Theorem 3 thus requires a bit of additional work, which is in Section B. Prior to this work, the best hardness ratio known for SDR with these parameters was $0.2/\ell$ [8,10] (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).

4 Overview of the Proof of Theorem 1

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, i.e., the ancestor robustness guarantee applies to the footprint of S on level b.*

4.1 Summary of the SDR Proof from Fisch [8]

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. There are just not enough levels and not enough degree for exponential growth to do its job, especially when pebbles slow down this growth.

The proof in [8] 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 predecesor-robustness for depth-robustenss):

1. **Expansion to get a large footprint of a large set.** Assume S is a subset of the bottom layer and $wt(S) \geq \zeta$. Prove, using vertical edges and expansion arguments, that there exists a fertile level. At its core, the argument is relatively simple: the set 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 [8] is a messy function, and the proof uses its piecewise-linear approximation to reach $\pi = 0.8$. 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 [8] (see Section B.1).
2. **Predecessor robustness to get a single-sink footprint.** By the predecessor robustness property, the footprint of S on level b contains a single-sink subgraph T of size 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, 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).

4.2 Main Idea of the Improvement

Our proof that footprint size approaches $r\ell$ for any $r < 1$ 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 r on almost every level above b.

Applying Step 4 to T' instead of T implies that there are at least $(1-\zeta)\cdot n$ nodes at the bottom level whose footprint is of size almost $r(\ell - b)$.

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 6. 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 7, 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 set. Moreover, some black pebbles can be used to stop several fertile sets at once. In Section 8, 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 9.

5 Proof Notation and Basic Notions

We recap notation used above and introduce some new notation.

5.1 Graph, Weights, Gains

Given a set S of nodes, let weight $wt(S)$ denote $|S|/n$. 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. The layers 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.[5]

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

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 $[\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}]$ 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 $\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}$ are fixed). Note that $\alpha_\delta^{\mathsf{min}} < \alpha_g < \alpha_\delta^{\mathsf{max}}$.

[5] 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 3.

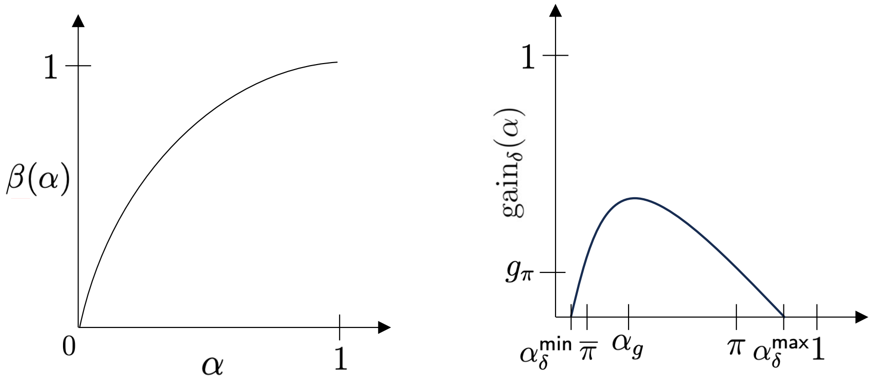

Fig. 1. Generic β and $gain_\delta$ (see Figure 2 in Section A for the Chung expander).

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

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 $\beta\text{count}_x(y) = \min\{i : \beta_\delta^i(x) \geq y\}$.*

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 23; 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 23, 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 used in SDR (see Section A for details), expansion is rapid all the way from $x = \delta$ to $y = 1 - 2\delta$, taking fewer than $2 \log 1/\delta$ steps, as we prove in Claim 10. Moreover, simple numerical computations for the Chung expander show $\beta\text{count}_\delta(1 - 2\delta) < 10$ for $\delta \geq 0.001$.

5.2 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., $\rho_{i\ldots j} = \rho_{j\ldots i}$).

Let $\zeta = (1 - \varepsilon_{\mathsf{space}})/2$ (this choice is somewhat arbitrary, and we could 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 b, define $f_b(T)$ to be the weight of the unpebbled part of T. For $i \leq b$, 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 $j \leq i$, $f_j = 0$, because $\beta_\delta(0) = 0$. When T is clear from the context, we will write $f_i(b)$ or simply f_i instead of $f_i(T)$ to simplify notation.

Define the functions ϕ_T and ϕ_{f_b} as

$$\phi_T(\rho_b \ldots, \rho_1) = \phi_{f_b}(\rho_{b-1} \ldots, \rho_1) = f_b + \cdots + f_1$$

to provide lower bounds on the total footprint of T.

We note that f and ϕ obey intuitive monotonicity constraints.

Claim 2. *For every $j \geq i$, f_i is monotone nonincreasing as a function of ρ_j and monotone nondecreasing as a function of $wt(T)$. The function ϕ is monotone: if $f'_b \geq f_b$ and $\rho'_i \leq \rho_i$ for each i, then $\phi_{f'_b}(\rho'_{b-1}, \ldots, \rho'_1) \geq \phi_{f_b}(\rho_{b-1}, \ldots, \rho_1)$. Moreover, adding a level at the end cannot decrease ϕ: $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1, \rho_0) \geq \phi_{f_b}(\rho_{b-1}, \ldots, \rho_1)$ 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_b to f'_b or from ρ to ρ'. The third sentence follows from nonnegativity of f. □

5.3 Basic Facts about Measuring Footprints

The following simple claim will turn out surprisingly useful in understanding 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.

$$f_m = \max\left(0, f_b + \mathit{gain}_\delta(f_b) + \cdots + \mathit{gain}_\delta(f_{m+1}) - \rho_{m\ldots b-1}\right)\,.$$

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

$$\beta_\delta(f_m) - \rho_{m-1} = f_m + \mathit{gain}_\delta(f_m) - \rho_{m-1} = f_b + \sum_{i=m}^{b} \mathit{gain}_\delta(f_i) - \rho_{m-1\ldots b-1}\,.$$

□

We generally will be interested in footprints that grow. The following claim allows us to rule out some situations in which a footprint decreases even without any black pebbles. This decrease can occur when the footprint 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 the footprint. The claim shows, in part, that if the footprint 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. *Therefore, 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 23), 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, $f_{i-1} > f_i$. □

6 Upperbounding the Number of Infertile Levels

Recall the definition of fertile (Definition 3) and constraints on ζ (Section 5.2). The main idea for bounding the number of infertile levels is the following. Take a subset S of nodes on the bottom level ℓ of weight ζ. Suppose some level b is infertile, which means the footprint weight $f_b(S) < \pi$. Consider two cases:

- If $f_b(S) \geq \overline{\pi}$, then the gain $gain_\delta(f_b(S))$ of level b is at least g_π, so there have to be enough pebbles so that the next infertile level above b can overcome this gain, per Claim 3. Thus, every infertile level (except the lowest) costs at least g_π in black pebbles. Note that this argument does not say when the g_π black pebbles must be placed, as long as they are above the infertile level.
- If $f_b(S) < \overline{\pi}$, most of the black pebbles must be at level b or below, per Claim 3, because $\overline{\pi}$ is small. The footprint will grow above b until it gets to π, and there are not many pebbles left to stop this growth, so as soon as the footprint reaches π, the remaining levels above will be fertile.

These cases essentially correspond to two possible adversarial strategies for placing pebbles: either keep every infertile level just below π and spend $g_\pi(\pi)$ black pebbles to keep it infertile one level up, or spend all the black pebbles at once to get a very small footprint, which will remain infertile for a few levels of growth. 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 infertile levels are nearly tight, as we further discuss below.

As a result, we obtain the following theorem.

Theorem 4. *Assume $g_\pi > 0$ and $\alpha_\delta^{\mathsf{min}} + \rho < \zeta_\delta < \alpha_\delta^{\mathsf{max}}$. 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. As we explain following Lemmas 1 and 2, the bound in this theorem is tight up to 1 level as long $\zeta_\delta \geq \pi$; else it is a slight overestimate.

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

Claim 5. *Assume $\alpha_\delta^{\mathsf{min}} + \rho < \zeta_\delta < \alpha_\delta^{\mathsf{max}}$. Then*

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

and $gain_\delta(f_m) > 0$.

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

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

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

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

Theorem 4 now follows from Lemmas 1 and 2 below. Note that the first argument to the max is greater than the second when $\zeta_\delta - \rho \geq \overline{\pi}$ by Claim 1. □

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

Lemma 1. *Assume $g_\pi > 0$ and $\alpha_\delta^{\mathsf{min}} + \rho < \zeta_\delta < \alpha_\delta^{\mathsf{max}}$. Let $m < \ell$ be some level. Assume there are $k > 0$ infertile levels below m, and assume that for all $i > m$, $f_i(S) \geq \overline{\pi}$. Then*

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

and thus the total number of infertile levels 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 first infertile 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. The footprint of every infertile level is at most π, so the footprint of every infertile level below m is between π and $\overline{\pi}$, so its gain is least g_π. The gains of other levels are positive by Claim 5. Let m' be the last highest infertile level below level m. Because it's infertile, $\pi > f_{m'}$, so by Claim 5

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

Rearranging the terms concludes the proof. □

6.2 The more complex case: small footprints

Lemma 2. *Assume $g_\pi > 0$ and $\alpha_\delta^{\mathsf{min}} + \rho < \zeta_\delta < \alpha_\delta^{\mathsf{max}}$. Assume for some level i, $f_i(S) < \overline{\pi}$. Then the number of infertile levels 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 $\beta\mathrm{count}_{\zeta_\delta - \rho}(\pi) - 1$ infertile levels.

Proof. 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 on 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 $\beta(\rho_{m+2}) \geq \pi$. But because m is infertile, we know

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

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

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

Recalling that $\beta(\rho_{m+2}) \geq \pi$, 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. □

7 Lowerbounding Footprints of Fertile Levels

In this section, we switch from thinking 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.

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 the total footprint, because growth stops earlier. It turns out that this intuition is not true in general, because the footprint on the higher of the two levels may grow slightly as the pebble moves down, which will cause the footprints in the levels above it to also grow, compensating for the reduction. For example, for the parameters of Section B, $\phi_{0.2}(0, 0.7, 0) \approx 1.007$, 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 5. *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 the total footprint*

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

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 footprint 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 $\rho_m, \dots, \rho_{b-1}$, the following holds as long as $gain_\delta(f_b) > 0$.*

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

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

Claim 6. *Suppose $f_i \leq \alpha_g$. 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 $x \geq 0$. 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 $f_{i-1}, \dots, f_1$ will increase, and thus ϕ will not increase. □

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

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

Denote the per-level footprint bounds 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 $g \geq h$. Because Claim 6 covers the case of $f_{m+1} = g_{m+1} \leq \alpha_g$, 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 $g_i \geq h_i$, 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 $g \geq h$. 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 $1 \leq i < m$, $f_{m-i} - g_{m-i} = \beta_\delta^i(f_m) - \beta_\delta^i(f_m - \rho_m)$, 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 $1 \leq i < m+1$, $f_{m-i+1} - h_{m-i+1} = \beta_\delta^i(f_{m+1}) - \beta_\delta^i(f_{m+1} - \rho_m)$. Thus,

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

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

Thus, to prove that $f - g < f - h$, it suffices to prove that $\beta_\delta^i(f_m) - \beta_\delta^i(f_m - \rho_m) \leq \beta_\delta^i(f_{m+1}) - \beta_\delta^i(f_{m+1} - \rho_m)$. 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 $\alpha_g > \alpha_\delta^{\mathsf{min}}$). Note also that β_δ^i, as a self-composition of a concave increasing function, is concave by repeated application of Claim 25. 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 26 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 C.

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 $f_{m+1} > \alpha_\delta^{\mathsf{min}}$, 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 $1 \leq i \leq m$, $g_i \geq h_{i+1}$.*

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, $g_m = \max(0, f_m - \rho_m) \geq \max(0, f_{m+1} - \rho_m) = h_{m+1}$.

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, $g = \sum_{i=1}^{b} g_i = \sum_{i=m+2}^{b} g_i + g_{m+1} + \sum_{i=1}^{m} g_i \geq \sum_{i=m+2}^{b} h_i + 0 + \sum_{i=1}^{m} h_{i+1} = h - h_1$. If any of the g_i values is ever 0, then $g_1 = 0$, so by the above claim $h_2 = 0$ (since $h_2 \leq g_1$), 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. □

8 Upperbounding the Number of Fertile Levels that Stop Growing

Thanks to Theorem 5, 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 the 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. We will consider the growth stopped if a footprint on some level becomes less than α_π. (While this will give a slightly suboptimal bound, because such a footprint may yet recover, we are only slightly undercounting the cost to the adversary: note that the footprint can be dropped to 0 with $\beta_\delta(\alpha_\pi)$ black pebbles, while to get a footprint 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. *Let T be an unpebbled set of weight at least α_π in level b. We will say that T (or simply level b) is* viable *for k levels if $f_{b-i}(T) \geq \alpha_\pi$ 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 T is* extinguished *after k levels if it is viable for k levels and $f_{b-k}(T) < \alpha_\pi$.*

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 6. *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 $\sigma > \alpha_\delta^{\mathsf{min}}$.*

Assume

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

Then there is a fertile level $b \geq m$ with 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; 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 6 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 $g_{\alpha_\pi} \geq 0$. Assume a subset T of level b is viable for k levels. Then the total of first k gains satisfies*

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

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 spending 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 (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 $f_m > \alpha_g > \alpha_\delta^{\mathsf{min}}$) 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 25, the latter because it's a line), and thus their sum is concave by 24, and thus the minimum is reached at the extrema of f_j by Claim 23. 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 set.

Corollary 1. *Assume $g_{\alpha_\pi} \geq 0$. Assume a subset T of 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 sets that have not been extinguished. We cannot bound the number of pebbles spent on such sets, 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 a subset T of 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 sets 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 set 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 $0 \leq m < k_1$), $gain_\delta(\beta_\delta^m(\alpha_\pi)) < g_{\alpha_\pi}$, which, by Fact 1, means $\beta_\delta^m(\alpha_\pi) > \alpha_g$ (because $\beta_\delta^m(\alpha_\pi) \geq \alpha_\pi$ by Claim 4). Take the smallest such m. By Claim 8, $gain_\delta(\beta_\delta^{j_1}(\alpha_\pi)) \leq gain_\delta(\beta_\delta^{j_2}(\alpha_\pi)) < g_{\alpha_\pi}$ for any $j_1 \geq j_2 \geq m$. From this step, we derive two inequalities.

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

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

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

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

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

Proof (Proof of Theorem 6). The assumptions imply the that the βcount term is finite by Claim 1, because $g_{\alpha_\pi} \geq 0$ implies $\alpha_\pi > \alpha_\delta^{\mathsf{min}}$.

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

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

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

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

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 that stays viable until level m inclusive. Then, letting $k' = b - m + 1$

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

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

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

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

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

By the condition on k in Theorem 6, the right-hand side of this inequality at least σ. We can thus apply Theorem 5 to level m and substitute σ instead of $\beta_\delta(f_m(b)) - \rho_{1\ldots m-1}$ by monotonicity of β (the condition $f_m \geq \alpha_\pi$ is satisfied because f_m is viable). □

9 Finishing the Proof with Quantitative Details

Theorem 1. *Fix $\varepsilon_{\mathsf{space}} > 0$. For any constant $r < 1$, there is a setting of constants δ, ℓ_{pr}, and n in the SPR construction, and a constant m, such that for any number of layers ℓ the SPR construction has hardness ratio*

$$r_{\mathsf{hardness}} \geq \frac{r(\ell - m)}{\ell}$$

(which approaches r as ℓ grows) and single-query catching probability $p_{\mathsf{hard}} \geq \varepsilon_{\mathsf{space}}/2$.

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

1. $\delta < gain(\pi)$ (so that $g_\pi > 0$ for Theorem 4)
2. $\alpha_\delta^{\mathsf{min}} < \varepsilon_{\mathsf{space}}/2 - \delta$ for Theorem 4 (this implies $\zeta_\delta - \rho > \alpha_\delta^{\mathsf{min}}$ because $\rho = 1 - \epsilon$)
3. $\alpha_\delta^{\mathsf{max}} + \delta > 1 - \varepsilon_{\mathsf{space}}/2$ for Theorem 4
4. $\delta < gain(\alpha_\pi)$ (so that $g_{\alpha_\pi} > 0$ for Theorem 6)
5. $\alpha_\delta^{\mathsf{min}} < \alpha_\pi$ for Theorem 6
6. $\alpha_\delta^{\mathsf{max}} - \alpha_\delta^{\mathsf{min}} > \rho$ for Theorem 6
7. $\alpha_\delta^{\mathsf{max}} > r$

Fix some σ so that $\rho + \sigma < \alpha_\delta^{\mathsf{max}}$ and $\sigma > \alpha_\delta^{\mathsf{min}}$ for Theorem 6 (this is possible by condition 6 on δ above). Set m_1 to be the largest integer smaller than

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

and m_2 to be the largest integer no greater than

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

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 4.

Therefore, Theorem 6 applies, which means that among the lowest ℓ_{pr} levels, there is a fertile level with footprint weight at least $\alpha_\pi + \sum_{i=0}^{\ell - \ell_{\mathsf{pr}} - 1} \beta_\delta^i(\sigma)$. Let $m_3 = \beta\text{count}_\sigma(r)$ (it's bounded by the choice of σ, Condition 7 on δ, and Claim 1). Then the last $\ell - \ell_{\mathsf{pr}} - m_3$ terms in this sum are at least r, and thus this footprint is of size at least $rn(\ell - m_1 - m_2 - m_3)$. The entire graph is of size $n\ell$ and the result follow by the proof in Section 4.2. □

Acknowledgments

Thanks to Carla Ràfols 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. Thanks to Protocol Labs for funding this work and to Carla Ràfols for hosting me at Universitat Pompeu Fabra.

References

1. Abusalah, H., Alwen, J., Cohen, B., Khilko, D., Pietrzak, K., Reyzin, L.: Beyond hellman's time-memory trade-offs with applications to proofs of space. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part II. LNCS, vol. 10625, pp. 357–379. Springer, Heidelberg (Dec 2017). https://doi.org/10.1007/978-3-319-70697-9_13

2. Alwen, J., de Rezende, S.F., Nordström, J., Vinyals, M.: Cumulative space in black-white pebbling and resolution. In: Papadimitriou, C.H. (ed.) ITCS 2017. vol. 4266, pp. 38:1–38:21. LIPIcs, 67 (Jan 2017). https://doi.org/10.4230/LIPIcs.ITCS.2017.38
3. Ateniese, G., Bonacina, I., Faonio, A., Galesi, N.: Proofs of space: When space is of the essence. In: Abdalla, M., Prisco, R.D. (eds.) SCN 14. LNCS, vol. 8642, pp. 538–557. Springer, Heidelberg (Sep 2014). https://doi.org/10.1007/978-3-319-10879-7_31
4. Benet, J., Dalrymple, D., Greco, N.: Proof of replication (2017), https://filecoin.io/proof-of-replication.pdf
5. Chung, F.: On concentrators, superconcentrators, generalizers, and nonblocking networks. Bell System Technical Journal **58**(8), 1765–1777 (1979)
6. Dziembowski, S., Faust, S., Kolmogorov, V., Pietrzak, K.: Proofs of space. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 585–605. Springer, Heidelberg (Aug 2015). https://doi.org/10.1007/978-3-662-48000-7_29
7. Fisch, B.: Tight proofs of space and replication. Cryptology ePrint Archive, Report 2018/702 (2018), https://eprint.iacr.org/2018/702
8. Fisch, B.: Tight proofs of space and replication. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part II. LNCS, vol. 11477, pp. 324–348. Springer, Heidelberg (May 2019). https://doi.org/10.1007/978-3-030-17656-3_12
9. Fisch, B., Bonneau, J., Greco, N., Benet, J.: Scaling proof-of-replication for filecoin mining (2018), https://research.protocol.ai/publications/scaling-proof-of-replication-for-filecoin-mining/fisch2018.pdf
10. Giacomelli, I., Nizzardo, L.: Filecoin proof of useful space — technical report. Tech. rep., CryptoNet — Protocol Labs (2023), https://drive.google.com/file/d/1not0bdkPT1BCztgspIpzSUAzWSrM8h81/view
11. Labs, P.: Filecoin spec. Algorithms. Stacked DRG PoRep (2023), https://spec.filecoin.io/algorithms/sdr/
12. Pietrzak, K.: Proofs of catalytic space. In: Blum, A. (ed.) ITCS 2019. vol. 124, pp. 59:1–59:25. LIPIcs (Jan 2019). https://doi.org/10.4230/LIPIcs.ITCS.2019.59
13. Ren, L., Devadas, S.: Proof of space from stacked expanders. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part I. LNCS, vol. 9985, pp. 262–285. Springer, Heidelberg (Oct / Nov 2016). https://doi.org/10.1007/978-3-662-53641-4_11

A Proof of Theorem 2

The Chung bipartite expander of degree d [5] 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 $\lfloor i/n \rfloor$), choose a random permutation p, and draw an edge from port i on the top part to port $p(i)$ on the bottom part.

As shown by Ren and Devadas [13, Theorem 1], 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

$$\frac{1}{n} 2^{n \cdot (\mathsf{H_b}(x) + \mathsf{H_b}(y) + d \cdot (y \mathsf{H_b}(x/y) - \mathsf{H_b}(x)))} ,$$

where $x = u/n$, $y = u/n$, and $\mathsf{H_b}$ is the binary entropy function.

The work of [13] used expansion at one point α, whereas we want it to work almost the entire $[0, 1]$ — specifically, on the interval $[\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}]$.

Given a security parameter λ (so that a random permutation satisfies the expansion condition with probability $1 - 2^{-\lambda}$), take $\varepsilon_{\mathsf{chung}} = \lambda/n$ and define

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

We need to make sure it is well defined (i.e., the set of y values is non-empty) on the entire interval $[\alpha_\delta^{\mathsf{min}}, \alpha_\delta^{\mathsf{max}}]$; this may require raising n in order to lower $\varepsilon_{\mathsf{chung}}$, as the right-hand side of the inequality gets closer to 0 for α closer to 0 and to 1. 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 $2^{-n\varepsilon_{\mathsf{chung}}} = 2^{-\lambda}$.

We will be using $d = 8$. Because we will be working with large graphs, we can have very small $\varepsilon_{\mathsf{chung}}$ (which helps make expansion faster even for small α) and still maintain a security parameter λ. For our numerical computations in Section B, we will use $\varepsilon_{\mathsf{chung}} = 2^{-20}$ and $n = 2^{30}$, which gives extremely high assurance that a random graph is an expander with expansion β: namely, security parameter $\lambda = 1024$. (We note that the computations barely change — by less than 1/1000 in the β values — even if we use a much higher $\varepsilon_{\mathsf{chung}}$ of 2^{-10}.)

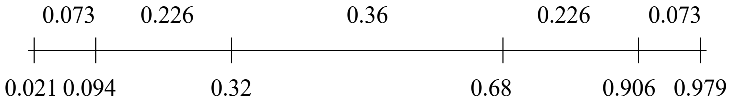

Fig. 2. Expansion speed for the Chung expander: $\beta(0.021) \approx 0.094$, $\beta(0.094) \approx 0.32$, etc. The top numbers show the gain.

Fact 2 *For a random Chung expander of degree 8 and sufficiently large n, with high probability*

1. $\alpha_g \approx 0.32$
2. *gain*$(\alpha_g) \approx 0.36$
3. *Expansion is rapid for small sets, with* $\beta(\alpha) > 3\alpha$ *for any constant* $0 < \alpha \leq 0.14$
4. *Expansion quickly bridges small-to-big gap:* $\beta(\beta((\beta(0.14) - 0.14) - 0.14) - 0.14 > 0.58$*, i.e.,* $\beta_\delta^3(0.14) > 0.58$ *for* $\delta \leq 0.14$
5. *Expansion quickly reaches almost everything for big sets:* $\beta(1 - 3\gamma) > 1 - \gamma$ *for any constant* $0 < \gamma \leq 0.14$

These facts imply that for any $\delta \le 0.14$, if there are no black pebbles, expansion can rapidly get the per layer footprint from δ to $1-2\delta$, as shown in the following claim.

Claim 10. *For the degree-8 Chung expander and $\delta \le 0.14$,*

$$\beta\text{count}_\delta(1-2\delta) \le 4 + 2\log_2(0.14/\delta) < -1.6 + 2\log(1/\delta)\,.$$

Proof. Let $k = \lceil \log_2(0.14/\delta) \rceil$. By Fact 2.3,

$$\beta_\delta(\alpha) > 3\alpha - \delta \ge 2\alpha$$

for $\delta \le \alpha \le 0.14$, so, by monotonicity of β (Condition 1), $\beta_\delta^k(\delta) > \max(2 \cdot 0.14, \delta \cdot 2^k) \ge \max(0.28, \delta \cdot 2^{\log_2(0.14/\delta)}) = 0.14$. By Fact 2.4 and monotonicity of β, $\beta_\delta^{k+3}(\delta) > 0.58$.

Now it's easier to proceed from the tail end. Note that $\beta_\delta^{-1}(x) = \beta^{-1}(x+\delta)$, so

$$\beta_\delta^{-1}(1-2\delta) = \beta^{-1}(1-\delta) = 1 - \beta(\delta) < 1 - 3\delta$$

by 2.5. And for $3\delta \le x \le 0.21$, we have $x - \delta \ge 2x/3$, so

$$\beta_\delta^{-1}(1-x) = \beta^{-1}(1-x+\delta) \le \beta^{-1}(1-2x/3) < 1-2x,$$

where the last step follows from $2x/3 \le 0.14$. Thus, by monotonicity of β, $\beta^{-(k+1)}(1-2\delta) = \min(1-2\cdot 0.21, 1-3\cdot 2^k\delta) \le \min(0.58, 1-3\cdot 2^{\log_2(0.14/\delta)}) = 0.58 < \beta_\delta^{k+3}(\delta)$. Applying β_δ^{k+1} to both sides gives the result of Claim 10 by monotonicity of β. □

Theorem 2. *For SPR instantiated with the degree-8 Chung expander, m in Theorem 1 is at most linear in*

$$\frac{1}{\beta(\alpha_\pi) - \alpha_\pi} + \frac{1}{\beta(\pi) - \pi}$$

plus an amount that is logarithmic in the inverses of $\varepsilon_{\mathsf{space}}, \alpha_\pi, 1-\pi$ and $1-r$.

Proof. Set

$$\delta < \min\left(\frac{gain(\pi)}{2}, \frac{gain(\alpha_\pi)}{2}, \frac{\varepsilon_{\mathsf{space}}}{4}, \frac{1-\pi}{2}, \alpha_\pi, 1-r, 0.14\right)$$

and $\sigma = \delta$. Use the proof of Theorem 1 from Section 9, observing that this setting of δ makes every βcount subscript at least δ and input at most $1-2\delta$, so by Claim 10, every βcount term is logarithmic in δ. □

B Proof of Theorem 3

Recall the parameters of Theorem 3, which match SDR as proposed by [8] and deployed by Filecoin [11]:

- a degree-8 Chung expander as described in Section A
- $\varepsilon_{\mathsf{space}} = 0.2$
- $\pi = 0.8$
- $\delta = 0.0378$ (using 180 challenges during initalization, with statistical soundness error 2^{-10})
- $\alpha_\pi = 0.2$ with the depth-robust graph of degree 6 (because proving practical depth robustness parameters is beyond our current state of knowledge, α_π is assumed, rather than proven [9])
- $n = 2^{30}$
- As before, we set $\zeta = 1 - \varepsilon_{\mathsf{space}}/2 = 0.9$ and thus $\zeta_\delta = 0.8622$.

In this section we use numerical estimates for values of β on specific inputs. We obtain these estimates via a simple implementation of the formula 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
- $0.3599 < gain(\alpha_g) < 0.3600$

In the rest of this section we proving Theorem 3, which we restate here.

Theorem 3. *Suppose SPR is instantiated with the degree-8 Chung expander, a predecessor-robust graph with $\pi = 0.8$ and $\alpha_\pi = 0.2$, $\ell_{\mathsf{pr}} = 8$, 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%.

For the rest of this section, let S be a set of weight ζ on level ℓ. We will use α_i to denote $f_i(S)$ (so as to easily distinguish it from $f_i(T)$).

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 4 (see Corollary 3), $m_2 = 3 = \beta\text{count}_{0.2}(0.9)$ from Theorem 6 (thus $\ell_{\mathsf{pr}} = m_1 + m_2 = 10$), and $m_3 = 4 = \beta\text{count}_{0.1}(0.92) = 4$. 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 $\ell = 11$.

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

- The proof Theorem 1 separately counts, and skips, the maximum number of infertile levels (Theorem 4) and the maximum number of fertile levels that are not going to grow (Theorem 6). 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 6 means fewer levels for Theorem 4 and vice versa.
- The proof of Theorem 6 ignores fertile levels whose footprints dip below α_π before rebounding and growing.
- The proof Theorem 1 ignores the footprint of a growing fertile level until it reaches weight r.

B.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 6 to the specific parameters of Theorem 3. However, Claims 14, 15, and 16 are new, and address the relationship between footprints, pebbles spent, and the number of fertile levels.

Claim 11. *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 4—namely, the one given by Lemma 1—applies and we have the following corollary to Theorem 4. Note that we get at most 7 for the number of infertile levels (we argue that this is tight in Section B.4), while the best previously known bound was 10 [8,10].

Corollary 3. *For the parameter settings in Theorem 3, 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 12. *Let m be the highest infertile level (assuming one exists).*

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

Proof. By Claim 5, $\sum_{i>m} gain_\delta(\alpha_i) \leq \alpha_m - \zeta + \rho + \delta$. 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 11 and level m is fertile so $\alpha_i < \pi$. □

Claim 13. *If i is the lowest fertile level, then $gain_\delta(\alpha_i) \geq 0.0313$.*

Proof. If $i = \ell$, then $\alpha_i \leq \zeta_\delta = 0.8622$. 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 $\alpha_i \geq \pi$ 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 4 leaves no levels for Theorem 6.

Claim 14. *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 $gain_\delta(\alpha_m) \leq 0.1501$. Indeed, this is automatically true for fertile m, because for a fertile m, $gain_\delta(\alpha_m) \leq g_\pi < 0.1114$. 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 $\sum_{i=m'}^{\ell} gain_\delta(\alpha_i) > gain_\delta(\alpha_b) + 0.1501 + 6 \cdot g_\pi > 0.0313 + 0.1501 + 6 \cdot 0.1113 = 0.8492$ (by Claim 11 and 13), which contradicts Claim 12.

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

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

so we must have $\alpha_m > 0.7418$.

Note that $\beta_\delta(\alpha_{m+1}) < \alpha_\delta^{\mathsf{max}}$ (else $gain_\delta(\beta_\delta(\alpha_{m+1})) \leq 0$, 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 15. *If there is m with $\alpha_m \leq 0.5015$, 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 11

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

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 $\alpha_m \leq 0.2023$.

Suppose $\alpha_m > 0.2023$. By Claim 23, because *gain* is concave (Fact 1), we know $\mathit{gain}_\delta(\alpha_m) \leq \min(\mathit{gain}_\delta(0.2023), \mathit{gain}_\delta(0.5015)) > 0.2804$ To have $\alpha_m \leq 0.5015$, 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

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

so $\mathit{gain}_\delta(\alpha_{m-1}) \geq \min(\mathit{gain}_\delta(0.0567), \mathit{gain}_\delta(0.7820)) > 0.1236$. 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 11

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

which contradicts Claim 12. This concludes the first case.

Now consider the second case: suppose $\alpha_m \leq 0.2023$. By Claims 5 and 11, because there are at least 5 infertile levels below the highest fertile level m', $\alpha_{m'} \geq 0.0622 + 5 \cdot 0.1113 = 0.6187 > 0.5015$. 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 $\alpha_{i+1} \leq 0.5015$ 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 16. *Suppose level $b \leq \ell - 6$ is infertile, and at least five levels below it are also infertile. If there is a level $m > b$ with $\rho_m > 0.2925$, then the weight $\rho_{1\ldots b-1}$ of black pebbles above b is at most* 0.0607.

Proof. 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, 11, and 13

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

Level m has $\alpha_m = \beta_\delta(\alpha_{m+1}) - 0.2925$; since $\beta_\delta(\alpha_{m+1}) < \alpha_\delta^{\mathsf{max}} < 0.9525$ (else $gain_\delta(\beta_\delta(\alpha_{m+1}) \leq 0$, which contradicts Claim 5) we have $\alpha_m < 0.66$. Since there are at least six infertile levels, by Claim 15, $\alpha_m > 0.5015$. 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 $\rho_{1 \ldots b-1} < \rho - 0.7393 = 0.0607$. □

B.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 Section B.3. The specific situations addressed in this section are:

- When $b \geq 4$ and $\rho_{b-1\ldots 1} \leq 0.07$ (Claim 17)
- When $b \geq 5$ and $\rho_{b-1\ldots 1} \leq 0.30$ (Claim 18)
- When $b \geq 6$ and $\rho_{b-1\ldots 1} \leq 0.44$ (Claim 19)
- When $b \geq 8$ and T is viable for at least three levels (Claim 20)
- When $b \geq 8$ and $\rho_{b-1} + \rho_{b-2} \leq 0.36$ and $\rho_{b-1\ldots 1} \leq 0.8$ (Claim 21)
- When $b \geq 8$ and $\rho_{b-1} \leq 0.1525$, $\rho_{b-1} + \rho_{b-2} \leq 0.73$, and $\rho_{b-1\ldots 1} \leq 0.8$ (Claim 22)

The first four of these claims simply apply the results of Sections 7 and 8 to the specific parameters of Theorem 3. 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 17. *Suppose $b \geq 4$, $f_b = 0.2$, and the total weight $\rho_{b-1\ldots 1}$ 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 5, 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 18. *Suppose $b \geq 5$, $f_b = 0.2$, and the total weight $\rho_{b-1\ldots 1}$ 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 5, 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 19. *Suppose $b \geq 6$, $f_b = 0.2$, and the total weight $\rho_{b-1\ldots 1}$ 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 5, 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 20. *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, $\beta_\delta(f_{b-2}) - \rho_{b-3 \ldots 1} = \beta_\delta(f_{b-2}) + \rho_{\ell \ldots b-2} - \rho > 0.1037$. Thus, by Theorem 5. the footprint of T is at least $0.2 \cdot 2 + \phi_{0.2}(\rho_{b-3 \ldots 1}, \underbrace{0, \ldots, 0}_{b-4}) = 0.6 + \phi_{0.1037}(\underbrace{0, \ldots 0}_{b-4}) \geq 0.6 + 2.8 + 0.94 \cdot (b-8)$. □

Claim 21. *Suppose $b \geq 8$, $f_b = 0.2$, $\rho_{b-1} + \rho_{b-2} \leq 0.36$, and the total weight $\rho_{b-1 \ldots 1}$ 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 $\phi_{f_b}(\rho_{b-1}, \ldots, \rho_1) \geq \phi_{f_b}(\rho_{b-1}, \rho_{b-2}, \rho_{b-3 \ldots 1}, 0, \ldots, 0) \geq \phi_{f_b}(\rho_{b-1}, \rho_{b-2}, 0.8 - \rho_{b-1} - \rho_{b-2}, 0, \ldots, 0)$ (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 22. *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 $\rho_{b-1 \ldots 1}$ 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 $\phi_{0.0782}(\underbrace{0, \ldots 0}_{b-4}) > 2.59 + 0.94 \cdot (b - 8)$. □

B.3 Putting the Proof of Theorem 3 Together

We now prove Step 3.5 described in Section 4.2, which suffices for proving Theorem 3.

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 17 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 18 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 19 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 14, and thus, by simple induction, m_1 can never be extinguished, so we can apply Claim 20 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 highest 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 20 to bound the total footprint. Else, set $b = m_2$. We know that the weight of black pebbles above level m_2 is at most 0.0607 by Claim 16, so we can apply Claim 17 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 18 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 19 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 19 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 21 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 $\rho_{m_2} \leq 0.1525$ because $\rho_{m_2} = \beta(m_1) - m_2 \leq \alpha_\delta^{\mathsf{max}} - \pi < 0.1525$ (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 $\rho_{m_2} + \rho_{m_2 - 1} \leq 0.73$, then set $b = m_1$ and apply Claim 22 to bound the total footprint. Else, there are at most five infertile levels by Claim 15, because $\alpha_{m_2 - 1} \leq 1 - \rho_{m_2 - 1} - \delta \leq$

$1 - (0.73 - 0.1525) - \delta = 0.3847$. Thus, there is a fertile level m_3 such that $\ell - 7 \leq m_3 \leq m_2 - 1$, 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 17 applies to bound the total footprint.

This concludes the proof of Lemma 5. □

B.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 $\rho_{\ell-9\ldots1} = 0$. 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 3. That doesn't mean the result can't be improved—perhaps other proof frameworks than the one in Section 4 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.

C 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 $a \geq b$, $F(a) \geq F(b)$ (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 23. *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 24. *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 25. *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 26. *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 5.1.

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 24. 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. □