Explicit Transformations from Edge-Depth-Robust to Node-Depth-Robust Graphs with Improved Concrete Efficiency

Jeremiah Blocki[1] and Nathan Smearsoll[1]

Purdue University, West Lafayette, USA
jblocki@purdue.edu, nsmearso@purdue.edu

Abstract. Depth-robust directed acyclic graphs (DAGs) are an important combinatorial primitive in cryptography, with applications to memory-hard functions, proofs of space, and proofs of sequential work. These applications require node depth-robustness, yet constructing sparse graphs with strong concrete guarantees remains a longstanding challenge. By contrast, edge depth-robust graphs are easier to construct explicitly and achieve better parameters, motivating the problem of efficiently transforming edge-depth-robust graphs into node-depth-robust graphs. We give a new and substantially more efficient transformation from edge-depth-robust graphs to node-depth-robust graphs. We give a new and substantially more efficient transformation from edge-depth-robust graphs to node-depth-robust graphs. Prior work of Blocki and Cinkoske (ITCS 2021) gave a general transformation by replacing each node with an ST-robust graph, resulting in extraordinarily large overhead and relying on non-explicit components. In contrast, our transformation replaces each node with a single superconcentrator rather than an ST-robust graph, yielding an explicit construction whenever the input graph is explicit, with dramatically improved concrete efficiency.

Formally, given an (e, d)-edge-depth-robust graph G with N nodes and m edges, our transformation generates a graph $G' = \textsc{Transform}(G_N)$ with $N' = O(m)$ nodes, constant indegree, and $(e/3,\ 2d-1)$-node depth-robustness. We also show that the transformation preserves fractional depth-robustness. Moreover, when $m = \omega(N)$, we show that our transformation can amplify depth — overcoming a key limitation of prior work. In particular, for any parameter d', we can obtain a constant-indegree graph G' with $N' = O(m+d'N)$ nodes and $(e/3,\ (d-e)d')$-node depth-robustness, yielding asymptotically tighter results in the important setting where $m = \omega(N)$ and $e \geq d$ by setting $d' \sim m/N$. As an application, We provide a novel analysis of Schnitger's edge-depth-robust graph construction (FOCS 1983). We show that the graph G_n, which has $N = 2^{n+1} - 1$ nodes and $m \leq n2^n$ edges, is $(e, 2d-1)$-edge-depth robust for all $e + d \leq 2^n$. Applying our transformation yields an explicit DAG on N' nodes with maximum indegree 2 and depth-robustness parameters $e = \Omega(N'/\log N')$ and $d = \Omega(N')$, matching the best known asymptotic trade-offs—even compared with non-explicit constructions—while improving the best previously known concrete ed-product lower bound by a multiplicative factor of 6.6.

1 Introduction

A directed acyclic graph (DAG) $G = (V, E)$ is (e, d) depth-robust if for any subset $S \subseteq V$ of at most $|S| \leq e$ nodes the graph $G-S$ obtained by deleting nodes in S (and any incident edges) still contains a directed path of length d. Depth-robustness has proven to be a useful combinatorial property in cryptographic constructions of Memory-Hard Functions [6,2,4], Proofs of Sequential Work [21] and Proofs of Space [16,22]. A growing line of research has focused on constructing sparse graphs that are maximally depth-robust e.g., [17,25,4,3,5,8,9]. It is known that depth-robust graphs are both necessary and sufficient to construct side-channel resistant Memory Hard Functions (called data-independent Memory Hard Functions (iMHFs)) with high cumulative memory complexity [2,4,6]. In particular, any iMHF based on an (e, d)-depth robust graph has cumulative pebbling complexity at least ed [4] and any graph that is *not* (e, d)-depth-robust (such graphs are called (e, d)-reducible) admits a pebbling attack with cumulative cost at most $O\left(eN + \frac{N^2 d}{e}\right)$ [2].

Ideally, to construct an iMHF we would like to have a DAG with the following properties: (1) the DAG has constant indegree (ideally maximum indegree is two), and (2) the graph is (e, d)-depth robust for parameters e, d which maximize the product. Erdos et al. [17] constructed a (e, d)-depth-robust graph with maximum indegree $O(\log N)$ and asymptotically optimal depth-robustness parameters $e = \Omega(N)$, $d = \Omega(N)$ and $ed = \Omega(N^2)$. Alwen et al. [4] showed how to transform this graph into a (e, d)-depth robust graph with maximum indegree 2, $e = \Omega(N/\log N)$ and $d = \Omega(N)$. Asymptotically, this is optimal for graphs with constant indegree [2]. However, this combinatorial construction is too expensive for practice and the constant factors hidden by by the asymptotic notation are highly undesirable.

Constructing sparse depth robust graphs with (provably) good concrete parameters has proven to be challenging. [3] constructs a practical graph (distribution) DRSample with N nodes and indegree 2 and prove that (whp) a randomly sampled graph from the distribution is $(c_1 N/\log N, c_2 N)$-depth robust where $c_1 \approx 0.0002$ and $c_2 = 0.03$. Asymptotically, this is optimal[1] but the provable depth-robustness constants c_1 and c_2 are both much smaller than we would desire e.g., applying [4] the lower bound on cumulative memory complexity would be $c_1 c_2 N^2/\log N$ where $c_1 c_2 \approx 7.3 \times 10^{-6}$.

Furthermore, most candidate constructions of depth-robust graphs are defined using a randomized distribution e.g., [17,3,7,10]. Even if we can show that a randomly sampled graph is (e, d)-depth robust with high probability there is not an absolute guarantee that the sampled graph is actually depth-robust. Testing if a graph is actually (e, d)-depth robust is computationally intractable [15,13]. This raises potential security concerns when the underlying graph is generated by an untrusted party. For example, a malicious party could generate a graph G that appears to be honestly sampled but secretly contains a small trapdoor

[1] If $e = \omega(N/\log N)$ then any graph G with constant indegree is (e, d)-reducible with $d = o(N)$[2].

depth-reducing set $S \subseteq V(G)$ such that $G - S$ contains no long directed paths. One notable exception is the explicit construction of Blocki and Lee [9] which achieves maximum indegree 2, $e = \Omega(N/\log N)$ and $d = \Omega(N)$. However, the authors acknowledge that the construction is "primarily of theoretical interest" because the hidden asymptotic constants are impractically large [9].

An alternative notion of depth-robustness is (e, d) edge-depth-robustness. A DAG $G = (V, E)$ is (e, d) edge-depth-robust if, after deleting any subset $S \subseteq E$ of at most $|S| \leq e$ *edges*, the remaining graph $G - S$ still contains a directed path of length d. In contrast to node-depth-robustness, it appears significantly easier to construct edge-depth-robust graphs that are explicit and achieve good concrete parameters [25,8]. For example, Schnitger [25] gave an explicit construction of a sparse edge-depth-robust graph G_n with $N = 2^{n+1} - 1$ nodes and $m = n2^n$ edges achieving strong tradeoffs between e and d. In particular, G_n guarantees depth depth $d = \Omega(N^{1-\epsilon})$ even when we can delete $e = \Omega(m)$ edges. We strengthen Schnitger's analysis in the important regime where $e \leq N$, showing that G_n is $(e, 2d - 1)$-edge depth-robust for all parameters satisfying $e + d \leq 2^n$. In particular, when $e = 2^{n-1}$ our analysis yields $d = 2^{n-1} = \Omega(N)$, whereas the original analysis of [25] only established $d = \Omega(N/\log N)$ in this regime.

Unfortunately, cryptographic applications generally require *node* depth-robustness instead of edge-depth robustness. In particular, Memory-Hard Functions [6,2,4], Proofs of Sequential Work [21] and Proofs of Space [16,22] all rely on node depth-robust graphs. This motivates the development of general transformations that convert edge-depth-robust graphs into node-depth-robust graphs.

Prior Transformations from Edge to Node Depth-Robustness Schnitger [25,24] gave a transformation $\textsc{Transform}_\epsilon$ using bipartite expander graphs to convert his edge-depth-robust graph G_n into a $(n2^n, 2^{(1-\epsilon)n})$-node-depth-robust graph $G_{n,\epsilon} = (V, E) = \textsc{Transform}_\epsilon(G_n)$ with $|V| = n2^{n+2}$ nodes and constant indegree. However, the transformation has two major limitations. First, $\textsc{Transform}_\epsilon$ can only be applied to $(\Omega(m), d)$-edge depth-robust graphs G with m edges. It does not provably preserve depth-robustness guarantees outside this narrow parameter regime. Second, the transformation $\textsc{Transform}_\epsilon$ depends on a fixed parameter ϵ, and the overhead grows rapidly as $\epsilon \to 0$, since the degree of the underlying bipartite expanders increases quickly with $1/\epsilon$. Consequently, the transformation preserves depth-robustness only at a single operating point (e, d), whereas the original graph G_n achieves strong edge-depth-robustness guarantees across a broad range of parameter pairs (e_i, d_i). In particular, for every $i \leq n$, the graph G_n is (e_i, d_i) edge depth robust, where $e_i = i \cdot 2^{n-1} = \frac{i|E_n|}{2n}$ and $d_i = \frac{2^n}{2\sum_{j=0}^{i}\binom{n}{j}}$.

Blocki and Cinkoske [8] gave the first *general* transformation from any (e, d)-edge depth robust graph G with N nodes and m edges to a regular (node) (e, d)-depth-robust graph G with $O(m)$ nodes and constant indegree. However, their transformation has several drawbacks. First, it relies on an *extremely* expensive combinatorial object called an ST-robust graph. Thus, while theoretically interesting, the construction is unlikely to be practical. Second, the transforma-

tion is not explicit, since all currently known constructions of ST-robust graphs either have superlinear size or rely on randomness. Finally, the analysis of [8] does not show how to achieve optimal depth parameters after transformation. In particular, their black-box analysis can only show that the transformed graph inherits path length $d \leq N$ from the original graph. Thus, if the original graph has $m = \omega(N)$ edges — as one would generally expect[2] — the resulting depth guarantee $d = o(m)$ is asymptotically smaller than the number of nodes in the transformed graph. Ideally, we would like the resulting node-depth-robust graph to achieve depth $d = \Omega(m)$, proportional to its total number of nodes.

Finally, [8] do not consider whether their transformation generically preserves fractional depth-robustness. Although their transformation likely does preserve this stronger property, no proof is provided.

Some stronger analyses of Memory-Hard Functions and Proofs of Space also utilize the notion of *fractional* depth-robustness, e.g., in the analysis of Argon2i [14] and in recent work on sustained-space and cumulative-memory tradeoffs for data-dependent memory-hard functions [11,12]. Intuitively, a graph is fractionally depth-robust if, after removing any sufficiently small set of vertices, there remain not just one but many vertices of large remaining depth. Thus, it is desirable to develop edge-to-node depth-robustness transformations that also preserve fractional depth-robustness. In fact, motivated in part by applications to Proofs of Space, [8] showed how to use ST-robust graphs to transform a node-depth-robust graph into a fractionally node-depth-robust graph. However, they did not analyze whether their original transformation from edge-depth-robustness to node-depth-robustness generically preserves fractional depth-robustness.

The fact that prior transformations yield useful—but still limited—constructions of depth-robust graphs motivates the main questions addressed in this work:

> Can we design an improved transformation TRANSFORM that efficiently maps arbitrary edge-depth-robust graphs into node-depth-robust graphs while preserving strong concrete guarantees?

More specifically, can such a transformation be explicit, thereby enabling efficient constructions of explicit node-depth-robust graphs?

1.1 Our Contributions

We give an improved transformation, TRANSFORM, which — given an (e, d) edge-depth robust graph $G_N = (V, E)$ with N nodes and m edges — uses superconcentrators to construct an $(e/3, 2d - 1)$ depth-robust graph G^*_N with $N' = O(m)$ nodes and constant indegree. Our transformation dramatically improves the concrete efficiency of the construction of Blocki and Cinkoske [8], which relied on an extremely expensive combinatorial object called an ST-robust graph. Existing constructions of ST-robust graphs require over 600 superconcentrators together with more than a thousand copies of another combinatorial

[2] If the original edge-depth-robust graph had constant indegree so that $m = O(N)$ then edge depth-robustness and node depth-robustness would already be asymptotically equivalent, eliminating the need for a transformation.

gadget called Grates [25] e.g., see discussion in [12, Appendix J]. By contrast, we show that each ST-robust graph can be replaced by a single superconcentrator.

This moves the approach from a primarily theoretical construction toward a potentially practical method for constructing node-depth-robust graphs with strong concrete parameters. We also note that explicit constructions of superconcentrators are known [1], whereas all currently known constructions of ST-robust graphs rely on randomness. Consequently, we obtain the first explicit transformation from edge-depth-robustness to node-depth-robustness. Since several known constructions of edge-depth-robust graphs are themselves explicit[25], our transformation immediately yields explicit constructions of node-depth-robust graphs.

We further improve on the analysis of Blocki and Cinkoske in two ways. First, we prove that our transformation preserves fractional depth-robustness, a stronger property used in advanced analyses of Memory-Hard Functions and Proofs of Space [14,11,12,8]. Second, we show that our transformation can preserve linear depth $d' = \Omega(N')$ in important parameter regimes where $e' = O(N'/\log N')$, and N' denotes the number of nodes *after* transformation.

Achieving linear depth after transformation is nontrivial when the original graph has $m = \omega(N)$ edges as we would expect[3]. In this regime, the transformed graph has $N' = \omega(N)$ nodes, so the original graph itself cannot contain paths of length $\Omega(N')$. Thus, the transformation must effectively *amplify* path length. For example, suppose we begin with an $(\Omega(N), \Omega(N))$-edge depth-robust graph with $m = \Theta(N \log N)$ edges. Our analysis shows that the transformed graph has $N' = \Theta(N \log N)$ nodes and is $(\Omega(N'/\log N'), \Omega(N'))$-depth robust. By contrast, prior black-box analysis would only yield $(\Omega(N'/\log N'), \Omega(N'/\log N'))$-depth robustness. This distinction is particularly important for constructing data-independent memory hard functions, since we can now obtain cumulative memory complexity lower bounds of $ed = \Omega(N'^2/\log N')$, matching the asymptotic guarantees of state-of-the-art constructions [3].

Finally, we show that there are concretely efficient constructions of $(\Omega(N), \Omega(N))$ edge depth robust graphs with $O(N \log N)$ edges. In particular, we provide a new analysis of the grates construction G_n of Schnitger [25] with $N = 2^{n+1} - 1$ nodes and $m = n2^n$ edges. While the original analysis established strong tradeoffs between e and d, it only implied edge depth-robustness with depth parameter $d = \Omega(N/\log N)$ in the regime $e = \Omega(N)$. By contrast, we prove that G_n is (e, d)-depth robust *for all* parameters e and d such that $e + d \leq 2^n$. In particular, this implies that G_n is $(N/4, N/4)$-edge depth robust. Combined with our transformation, this yields a concretely efficient construction of an indegree two (e, d)-node-depth-robust graph with $e = cN'/\log N'$, $d = c'N'$. This matches the best known asymptotic tradeoffs while substantially improving concrete efficiency over prior constructions [3].

We obtain a (non-explicit) construction of a (e, d)-node-depth-robust graph with N nodes, maximum indegree 2, and $ed \geq \frac{1}{7726.41}\frac{N^2}{\log N}$. This improves the best previously known concrete constant for the $ed = \Omega(N^2/\log N)$ regime by

[3] If every node has constant indegree so that $m = O(N)$ then edge/node depth-robustness are essentially equivalent concepts.

roughly a factor of 18 compared with DRSample [3]. We also obtain an explicit construction with $ed \geq \frac{1}{20736}\frac{N^2}{\log N}$, which still improves the corresponding DRSample constant by roughly a factor of 6.6. The ed product is a particularly important concrete measure for depth-robust graphs, especially in the analysis of memory-hard functions, because ed lower bounds the cumulative pebbling complexity of any (e, d)-depth-robust graph [4].

Moreover, because the underlying Grates construction G_n achieves strong edge-depth-robustness even in the regime $e = \Omega(N)$, the transformed graph also preserves meaningful node-depth-robustness guarantees when $e = \Omega(N')$. This parameter regime is particularly relevant in applications to Proofs of Sequential Work and in the analysis and construction of data-dependent memory-hard functions [8,11,12].

2 Preliminaries

Graph Notation Given a DAG $G = (V, E)$ and a node $v \in V$ we define $\texttt{parents}(v, G) = \{u : (u, v) \in E\}$ to be the set of all nodes u with an outgoing edge to v. The indegree of a node $v \in V$ is $\texttt{indeg}(G, v) \doteq |\texttt{parents}(v, G)|$. Similarly, the outdegree of a node $v \in V$ is $\texttt{outdeg}(G, v) \doteq |\{u : (v, u) \in E\}|$. Finally, it will be convenient to define the max-degree $\texttt{mdeg}(G, v) \doteq \max(\texttt{indeg}(G, v), \texttt{outdeg}(G, v))$ for a node v to be the maximum of the indegree and outdegree. When G is clear from context, we often drop G and write $\texttt{indeg}(v)$, $\texttt{mdeg}(v)$, or $\texttt{parents}(v)$. Similarly, we define the indegree/max-degree of a DAG G as $\texttt{indeg}(G) \doteq \max_{v \in V} \texttt{indeg}(G, v)$ (resp. $\texttt{mdeg}(G) = \max_{v \in V} \texttt{mdeg}(G, v)$).

Depth of a Graph: Given a directed path $P = v_1, v_2, \ldots, v_i$ in G such that $\forall_{1<j\leq i}, (v_{j-1}, v_j) \in E$, we say that the length of P, denoted $\texttt{length}(P)$, is the number of nodes in P, i.e. $\texttt{length}(P) = i$. We say that the depth of a vertex $v \in V$ is the length of the longest directed path ending at v, denoted $\texttt{depth}(v, G)$. The depth of a DAG G is the maximum depth of any vertex, $\texttt{depth}(G) \doteq \max_{v \in V} \texttt{depth}(v, G)$.

Reverse Graph: We define the reverse $\texttt{Rev}(G)$ graph, to be the graph obtained by reversing the direction of all edges in E, i.e., $\texttt{Rev}(G) \doteq (V, \texttt{Rev}(E))$ where $\texttt{Rev}(E) \doteq \{(v, u) : (u, v) \in E\}$. Observe that $\texttt{Rev}(\texttt{Rev}(G)) = G$ and that for all nodes $v \in V$ we have $\texttt{indeg}(G, v) = \texttt{outdeg}(\texttt{Rev}(G), v)$, $\texttt{outdeg}(G, v) = \texttt{indeg}(\texttt{Rev}(G), v)$ and $\texttt{mdeg}(G, v) = \texttt{mdeg}(\texttt{Rev}(G), v)$.

Removing Edges/Nodes: Given a DAG $G = (V, E)$ and a subset of vertices $S \subseteq V$, we define $G{-_V}S \doteq (V - S, \{(u, v) \in E \text{ s.t. } u \notin S \wedge v \notin S\}$ to be the subgraph obtained by removing all nodes in S and any incident edges. Similarly, given a DAG $G = (V, E)$ and a set of edges $S \subseteq V$, we define $G{-_E}S \doteq (V, E - S)$ to be the subgraph obtained by removing all the edges in S.

Node Depth Robustness: We say that a DAG G is (e, d)-depth reducible if there exists a set of vertices $S \subseteq V(G)$ of size $|S| = e$ so that $G{-_V}S$ contains no directed path of length d, i.e. $\texttt{depth}(G{-_V}S) < d$. If G is not (e, d)-depth reducible then we way G is (e, d)-depth robust, i.e. $\forall S \subseteq V(G)$ such that $|S| \leq e$ we have $\texttt{depth}(G{-_V}S) \geq d$.

Edge Depth Robustness We similarly say that a DAG G is (e, d)-*edge* depth reducible if there exists a set of *edges* $S \subseteq E(G)$ of size $|S| = e$ so that $G-_E S$ contains no directed path of length d, i.e. $\texttt{depth}(G-_E S) < d$, and that a DAG G is (e, d)-edge depth robust otherwise, i.e. $\forall S \subseteq E(G)$ with $|S| \leq e$ we have $\texttt{depth}(G-_E S) \geq d$.

Superconcentrators We refer to an n-superconcentrator as $\text{SC}(n)$. For a given superconcentrator $\text{SC}(n)$, we refer to the set of input and output vertices as $\text{IN}(\text{SC}(n))$ and $\text{OUT}(\text{SC}(n))$, respectively.

Definition 1 ([23,27]). *A n-superconcentrator is a DAG with disjoint subsets $\text{IN} \subset V(G)$ and $\text{OUT} \subset V(G)$ of size n such that, for every r in $1 \leq r \leq n$, for every subset $a \subseteq \text{IN}$ and subset $b \subseteq \text{OUT}$ with $|a| = |b| = r$, there exists a set of r vertex-disjoint directed paths connecting a to b.*

Pippenger [23] showed that there exist families of superconcentrators with linear size, i.e., superconcentrators $\text{SC}(n)$ having $|V(\text{SC}(n))| = O(n)$ vertices and $|E(\text{SC}(n))| = O(n)$ edges. A substantial line of work has focused on improving the concrete density $|E(\text{SC}(n))|/n$ of such constructions. Schöning [26] constructed a family of superconcentrators with density 28, later improved to density 25.3 by Kolmogorov and Rolínek [19]. Gabber and Galil [18] gave an explicit linear-size construction of superconcentrators, while Alon and Capalbo [1] later improved the density of explicit constructions to 44. Lev and Valiant [20] lower bounded the density of any super-concentrator as $5 - o(1)$. They also gave a recursive construction with density $3 \log_2 n$; in particular, this density is smaller than 25.3 whenever $n < 345$.

Finally, we note that any superconcentrator $\text{SC}(n)$ can be transformed into a superconcentrator with maximum indegree 2 by replacing each node v of indegree $\texttt{indeg}(\text{SC}(n), v)$ with a directed path of length $\texttt{indeg}(\text{SC}(n), v)$ and distributing the incoming edges along the path — see Lemma 12 in the appendix. While this transformation increases the density, it increases the number of edges by at most a factor of two. Moreover, if the original superconcentrator has m edges, then the transformed superconcentrator contains at most $m + n$ vertices. Thus, for our applications, the key parameter is not the density itself, but rather the total number of vertices in an indegree-2 superconcentrator, which is upper bounded by the density of the original construction. With this transformation, we have explicit constructions of indegree-2 superconcentrators with $45n$ nodes and randomized constructions with $26.3n$ nodes. Lev and Valiant [20] proved that any indegree-2 superconcentrator must contain at least $4n - o(n)$ vertices. We leave it as an interesting open question to determine the best achievable density for indegree-2 superconcentrators.

3 Technical Overview

3.1 Our Transformation

First, we give our transformation for any graph G_N (with N nodes). For each vertex $v \in V(G_N)$ of the graph, we select a $\texttt{mdeg}(G_N, v)$-superconcentrator

and replace node v with a fresh copy $\text{SC}_v = \text{SC}(\texttt{mdeg}(G, v))$ of this superconcentrator. We then connect every edge between two nodes to come from an output of the source node's superconcentrator and lead to an input of the destination node's superconcentrator to obtain the transformed graph $G_N^* = \textsc{Transform}(G_N, \text{SC}, I, O)$ — see Figure 1.

To describe the transformation more precisely it will be convenient to let $\text{IN}(\text{SC}_v)$ (resp. $\text{OUT}(\text{SC}_v)$) denote the set of input (resp. output) nodes in the superconcentrator SC_v corresponding to v. It will also be useful to fix injective functions $\texttt{in} : E(G_N) \to \bigcup_{v \in V(G_n)} \text{IN}(\text{SC}_v)$ and $\texttt{out} : E(G_N) \to \bigcup_{v \in V(G_N)} \text{OUT}(\text{SC}_v)$ such that for each directed edge $(u, v) \in E(G_N)$ from the original graph we have $\texttt{in}(u, v) \in \text{IN}(\text{SC}_v)$ and $\texttt{out}(u, v) \in \text{OUT}(\text{SC}_v)$. In particular, each edge (u, v) is mapped to a particular input (resp. output) node of SC_v (resp. SC_u) by $\texttt{in}$ (resp. $\texttt{out}$). Now we can define the graph transformation $\textsc{Transform}(G_N, \text{SC}, I, O) = G_N^* = (V_n^*, E_n^*)$ as follows:

$$V_n^* \doteq \left(\bigcup_{v \in V(G_N)} V(\text{SC}_v) \right) \quad E_n^* \doteq \left(\bigcup_{v \in V(G_N)} E(\text{SC}_v) \right) \cup \left(\bigcup_{e \in E(G_n)} (\texttt{out}(e), \texttt{in}(e)) \right) .$$

See Definition 3 for a fully formal definition of the transformation.

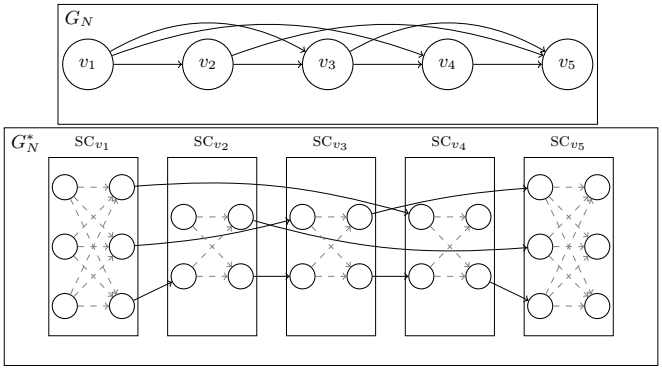


Fig. 1: Transformation from an edge-depth-robust graph G_N to a node-depth-robust graph $G_N^* = \textsc{Transform}(G_N, \text{SC}, I, O)$ using superconcentrators.

3.2 Proving Our Transformation is Depth Robust

We now wish to show that the resulting graph $G_N^* = \textsc{Transform}(G_N, \text{SC}, I, O)$ is node-Depth Robust as long as the original graph is edge-Depth Robust. To do so, we give a mapping from any subsetet $S^* \subseteq V(G_N^*)$ of deleted nodes from G_N^* to a corresponding set of edges $S \in E(G_N)$ of size at most $3|S^*|$ and we will argue that $\texttt{depth}(G_N^* - _V S^*) \geq 2 \cdot \texttt{depth}(G_N - _E S) - 1$. It will immediately follow that G_N^* is $(e/3, 2d - 1)$-node depth robust.

Definition 2. *Given any superconcentrator* SC_v *in* G^*_N *and a set* S^* *of nodes removed from* G^*_N*. We define the set* $S^*_v \doteq S^* \cap \text{SC}_v$ *as the set of all nodes removed from that superconcentrator and we define* $\gamma_v(S^*) = |S^*_v|/|\text{IN}(\text{SC}_v)|$*. When* S^* *is clear from context we will simply write* γ_v*.*

Then, we define two types of bad edges which must be in S to ensure a path in $G_N - _E S$ corresponds to a path in $\textsc{Transform}(G_N, \text{SC}, I, O) - _V S^*$. The first, and simpler, are "excluded" edges related to spoiled superconcentrators. Intuitively, a superconcentrator SC_v is spoiled if it has been so corrupted by S^* that we no longer make any guarantees about the corresponding node v in $G_N - _E S$. Formalized in Definitions 5 and 6, a superconcentrator is spoiled if the number of nodes removed from it is at least two thirds the size of its input — i.e. $|S^*_v| \geq \frac{2}{3}|\text{IN}(\text{SC}_v)|$. If SC_v is spoiled then we can also consider node $v \in V(G_N)$ to be spoiled along with all edges incident to v. Let $\textsc{ToExclude}(G_N, \text{SC}, I, O, S^*) \subseteq E(G_N)$ denote the set of spoiled edges in G_N. It takes $|S^*_v| \geq \frac{2}{3}|\text{IN}(\text{SC}_v)|$ nodes to spoil node v and spoiling node v will add at most $2 \cdot \texttt{mdeg}(v) \leq 2|\text{IN}(\text{SC}_v)|$ edges to $\textsc{ToExclude}(G_N, \text{SC}, I, O, S^*)$. Thus, we trivially can upper bound $|\textsc{ToExclude}(G_N, \text{SC}, I, O, S^*)| \leq 3|S^*|$, but this is not tight enough since we will need to add other edges to our set S. Observe that we can partition S^* into two sets $S^*_{spoiled} = \bigcup_{v \text{ spoiled}} S^*_v$ and $S^*_{unspoiled} = S^* \setminus S^*_{spoiled}$. A tighter upper bound is $|\textsc{ToExclude}(G_N, \text{SC}, I, O, S^*)| \leq 3|S^*_{spoiled}|$ — see Lemma 1. It will then suffice to argue that $|S \setminus \textsc{ToExclude}(G_N, \text{SC}, I, O, S^*)| \leq 3|S^*_{unspoiled}|$.

The more complicated types of bad edges are those which belong to non-spoiled superconcentrators, but are still "disrupted" by the removal of S^*. Before defining disrupted edges in G_N we first define three such sets of disrupted nodes in G^*_N: "destroyed" nodes $\textsc{Dest}(G^*_N, S^*)$, "internally disrespectful" nodes $\textsc{IntDisr}(G^*_N, S^*)$, and "externally disrespectful" nodes $\textsc{ExtDisr}(G^*_N, S^*)$. For each of these sets of nodes in G^*_N, there will a corresponding set of edges — $S_{\textsc{Dest}}(G^*_N, S^*)$, $S_{\textsc{IntDisr}}(G^*_N, S^*)$, and $S_{\textsc{ExtDisr}}(G^*_N, S^*)$ — in G_N which must be added to S.

Intuitively, the set of "destroyed" nodes $\textsc{Dest}(G^*_N, S^*)$ consists of all input nodes (for their corresponding superconcentrator) which are directly contained in S^*, and thus trivially cannot be used in a path in $G^*_N - _V S^*$. The corresponding set $S_{\textsc{Dest}}$ of "destroyed" edges $S_{\textsc{Dest}}$ contains all edges (u, v) in G_N whose corresponding endpoint node in G^*_N is destroyed. More formally:

$$\textsc{Dest}(G^*_N, S^*) = \bigcup_{v \in V(G_N)} \text{IN}\left(\text{SC}_v\right) \cap S^*_v$$

$$S_{\textsc{Dest}}(G^*_N, S^*) \doteq \{e \in E(G_N) \ : \ \texttt{in}(\texttt{e}) \in \textsc{Dest}(G^*_N, S^*)\}$$

See Definitions 7 and 10 for fully formal definitions of $\textsc{Dest}(G^*_N, S^*)$ and $S_{\textsc{Dest}}(G^*_N, S^*)$, respectively. It follows immediately that $|S_{\textsc{Dest}}(v, G^*_N, S^*)| \leq |S^*_v|$ for each node v where $S_{\textsc{Dest}}(v, G^*_N, S^*)$ denotes the set of destroyed edges incident to node v. Thus,

$$|S_{\textsc{Dest}}(G^*_N, S^*) \setminus \textsc{ToExclude}(G_N, \text{SC}, I, O, S^*)| \leq \sum_{v \text{ unspoiled}} |S^*_v| = |S^*_{unspoiled}|.$$

Next in order to define "internally disrespectful" and "externally disrespectful" nodes, we first define a useful function r, fully formalized in Definition 4, as follows:

$$r(v, G_N^*, S^*) \doteq \max\{t \ : \ |\{x \in \text{IN}(\text{SC}_v) \ : \ \texttt{depth}(x, G_N^* - _V S^*) \geq t\}| > |S_v^*|\}$$

Intuitively, the r value of a node $v \in G_N$ tells gives us the maximum depth t for which we can guarantee that sufficiently many the superconcentrator SC_v's inputs have are at depth at least t in $G_N^* - _V S^*$. By requiring that at least $|S_v^*|+1$ inputs have depth t we can invoke properties of the superconcentrator SC_v to ensure that all but $|S_v^*|$ outputs have depth at least $t+1$. Observe that as more nodes are removed from SC_v, i.e. $|S_v^*|$ is large, we require a larger fraction of the inputs to have depth r to preserve at least one path from such an input that is not disrupted. However, if the node v is not already spoiled, we can assume that this fraction does not exceed 2/3 .

Since every output node in a superconcentrator has a path from the highest-depth input node(s) in a superconcentrator, we might expect that in every superconcentrator SC_v the depth will increase, i.e. $\forall u \in \text{OUT}(\text{SC}_v), \texttt{depth}\,(u, G_N^* - _V S^*) \geq r(v, G_N^*, S^*) + 1$. However, due to the removal of nodes in S^* this may no longer be the case. Some output vertices may "disrespect" the expected increasing order of depths e.g., because that output node was removed or perhaps all paths from input nodes of depth at least $r(v, G_N^*, S^*)$ to this output node pass through a removed node. If $\texttt{depth}\,(u, G_N^* - _V S^*) < r(v, G_N^*, S^*) + 1$ for some output node $u \in \text{OUT}(\text{SC}_v)$ we call u "internally disrespectful" and we use $\text{IntDisr}(G_N^*, S^*)$ to denote the set of all such internally disrespectful nodes (Definition 8). We add an edge $(u, v) \in E(G_N)$ to the set S_{IntDisr} if $\texttt{out}(e) \in \text{IntDisr}(G_N^*, S^*)$ i.e., the corresponding edge $(\texttt{out}(e), \texttt{in}(e))$ in G_N^* begins at a node $\texttt{out}(e)$ which is internally disrespectful — see Definition 10. Formally, we have

$$\text{IntDisr}(G_N^*, S^*) \doteq \bigcup_{v \in V(G_N)} \{x \in \text{OUT}(\text{SC}_v) \ : \ \texttt{depth}\,(x, G_N^* - _V S^*) \leq r\,(v, G_N^*, S^*)\}$$

$$S_{\text{IntDisr}}(G_N^*, S^*) \doteq \{e \in E(G_N) \ : \ \texttt{out(e)} \in \text{IntDisr}(G_N^*, S^*)\}$$

To upper bound $|S_{\text{IntDisr}}(G_N^*, S^*)|$ we appeal to the properties of a superconcentrator. We claim that $|S_{\text{IntDisr}}(v, G_N^*, S^*)| \leq |S_v^*|$ for each node v where $S_{\text{IntDisr}}(v, G_N^*, S^*)$ denotes the set of internally disrespectful edges incident to node v — see Lemma 2. Thus, $|S_{\text{IntDisr}}(G_N^*, S^*) \backslash \text{ToExclude}(G_N, \text{SC}, I, O, S^*)| \leq \sum_{v \text{ unspoiled}} |S_v^*| = |S_{unspoiled}^*|$. To see why Lemma 2 holds, recall that, by definition of r, at least $|S_v^*| + 1$ input nodes have a depth *at least* $r(v, G_N^*, S^*)$. If $|S_{\text{IntDisr}}(v, G_N^*, S^*)| > |S_v^*|$ then we would have at least $|S_v^*| + 1$ node disjoint paths from our high depth input nodes (with depth at least $r(v, G_N^*, S^*)$) to these $|S_v^*| + 1$ internally disrespectful outputs and at least of one these paths would not contain a node from S_v^*. This leads to an immediate contradiction because the corresponding output node $o \in S_{\text{IntDisr}}(v, G_N^*, S^*)$ would now have depth at least $r(v, G_N^*, S^*) + 1$ and would therefore not be internally disrespectful.

Intuitively, if $(u, v) \in E(G_N)$ we expect that $r(v, G_N^*, S^*) \geq r(u, G_N^*, S^*) + 2$. However, this may not be the case for all edges. Let $e^* = (\texttt{out}(e), \texttt{in}(e)) \in$

$E(G^*_N)$ be the edge corresponding to $e = (u, v) \in E(G_n)$ in the original graph. If $r(v, G^*_N, S^*) \leq r(u, G^*_N, S^*) + 1$ then we say that $\texttt{in}(e)$ is externally disrespectful as long as (1) u, v are not spoiled, (2) $\texttt{out}(e)$ is not internally disrespectful and (3) $\texttt{in}(e)$ is not destroyed. In this case we would add $\texttt{in}(e)$ to the set $\textsc{ExtDisr}(G^*_N, S^*)$ of all externally disrespectful nodes and we would add (u, v) to the set $S_{\textsc{ExtDisr}} \doteq \{e \in E(G_N) : \texttt{in}(\texttt{e}) \in \textsc{ExtDisr}(G^*_N, S^*)\}$ of externally disrespectful edges — see Definitions 9 and 10.

We now need to bound the number $|S_{\textsc{ExtDisr}}(v, G^*_N, S^*)|$ of externally disrespectful nodes in a given superconcentrator where $S_{\textsc{ExtDisr}}(v, G^*_N, S^*) = \textsc{sc}_v \cap \textsc{ExtDisr}(G^*_N, S^*)$. Each externally disrespectful node $b \in S_{\textsc{ExtDisr}}(v, G^*_N, S^*)$ must have $\texttt{depth}(b, G^*_N - _V S^*)$ at least $r(u, G^*_n, S^*) + 2$. To see this note that there exists an edge $(a, b) \in E(G^*_N)$ such that $a \in \textsc{out}(\textsc{sc}_u)$ is not internally disrespectful meaning that $\texttt{depth}(a, G^*_N) \geq r(u, G^*_N, S^*)+1$ and $\texttt{depth}(b, G^*_N) \geq \texttt{depth}(a, G^*_N)+1$. By the definition of r, we cannot have more than $|S^*_v|$ nodes in $\textsc{in}(\textsc{sc}_v)$ with depth greater than $r(v, G^*_N, S^*)$. Thus, we can show that $|S_{\textsc{ExtDisr}}(v, G^*_N, S^*)| \leq |S^*_v|$ — see Lemma 3. It follows that $|S_{\textsc{ExtDisr}}(G^*_N, S^*) \backslash \textsc{ToExclude}(G_N, \text{SC}, I, O, S^*)| \leq \sum_{v \text{ unspoiled}} |S^*_v| = |S^*_{unspoiled}|$.

Now that we have defined and bound the size of all our sets of bad edges, $\textsc{ToExclude}(G_N, \text{SC}, I, O, S^*)$, $S_{\textsc{Dest}}(G^*_N, S^*)$, $S_{\textsc{IntDisr}}(G^*_N, S^*)$, and $S_{\textsc{ExtDisr}}(G^*_N, S^*)$, we use these sets, as well as the set of edges to/from spoiled superconcentrators $\textsc{ToExclude}(G^*_N, S^*)$ to define the set S of edges which must be removed from the original graph G_N.

$$S \doteq \textsc{ToExclude}(G_N, \text{SC}, I, O, S^*) \cup S_{\textsc{Dest}}(G^*_N, S^*) \cup S_{\textsc{IntDisr}}(G^*_N, S^*) \cup S_{\textsc{ExtDisr}}(G^*_N, S^*)$$

This is formalized in Definition 10. Further, since we know it follows from our earlier bounds that $|S \setminus \textsc{ToExclude}(G_N, \text{SC}, I, O, S^*)| \leq 3|S^*_{unspoiled}|$, and therefore that $|S| \leq 3|S^*|$. We formalize this proof in Lemma 4.

We now argue that removing S^* from G^*_N is similar to removing S from G_N i.e., $\texttt{depth}(G^*_N - _V S^*) \geq 2\texttt{depth}(G_N - _E S) - 1$. Intuitively, if we have a path $P(v_d) = v_1, v_2, \ldots, v_d$ in $G_N - _E S$ then we can argue that $r(v_i, G^*_N, S^*) > r(v_{i-1}, G^*_N, S^*) + 1$ for all $i \leq d$. This immediately implies that $r(v_d, G^*_N, S^*) \geq 2d-2$ — see Lemma 5. Briefly, we argue that if $r(v_i, G^*_N, S^*) \leq r(v_{i-1}, G^*_N, S^*)+ 1$ then the edge (u, v) would have necessarily been added to the set S and deleted. In u or v were spoiled then $(u, v) \in S$. Similarly, if $\texttt{in}((u, v)) \in \textsc{Dest}(G^*_N, S^*)$ or $\texttt{out}((u, v)) \in \textsc{IntDisr}(G^*_N, S^*)$ then $(u, v) \in S$ will be destroyed. Finally, if u, v are not spoiled $\texttt{in}((u, v)) \notin \textsc{Dest}(G^*_N, S^*)$, $\texttt{out}((u, v)) \notin \textsc{IntDisr}(G^*_N, S^*)$ and $r(v_i, G^*_N, S^*) \leq r(v_{i-1}, G^*_N, S^*) + 1$ then it necessarily follows that $\texttt{in}((u, v)) \in \textsc{ExtDisr}(G^*_N, S^*)$ meaning that (u, v) will be deleted.

We have proven that our transformed graph $G^*_N = \textsc{Transform}(G_N, \text{SC}, I, O)$ is node-depth robust.

Theorem 1 (Informal, See Theorem 5). *Given a graph G_N which is (e, d)-*edge *depth robust,* $\textsc{Transform}(G_N, \text{SC}, I, O)$ *is $(e/3, 2d-1)$-*node *depth robust for all $d > 1$.*

3.3 Extending our Transformation to Preserve Linear Depth

We have now shown that our transformation converts any (e, d) edge-depth robust graph into an $(e/3, 2d-1)$ node-depth robust graph. However, as long as $N = |V(G_N)| = o(|E(G_N)|)$, as we would expect in any useful construction of an edge-depth robust graph, we will necessarily have $d = o(|V(G^*_N)|)$ where G^*_N is our node depth-robust graph. In many applications we would like to construct graphs that are (e, d)-node depth robust when $d = \Omega(N)$ e.g., [3,4] construct (e, d)-node depth robust graphs with $e = \Omega(N/\log N)$, $d = \Omega(N)$ and constant indegree. At first glance it seems impossible to achieve a similar result using our transformation.

We address this gap by introducing the notion of a "d'-perfectly connected" superconcentrator — see Definition 13. Intuitively, a superconcentrator $\text{SC}(n)$ is d'-perfectly connected if for all inputs $u \in \text{IN}(\text{SC}(n))$ and all outputs $v \in \text{OUT}(\text{SC}(n))$ there is a path of length d' connecting u to v. These long paths will not necessarily be preserved once we start deleting nodes from $\text{SC}(n)$, but will be particularly useful in the special case that no node in $\text{SC}(n)$ is deleted. If we start with a family SC^{depth} of "d'-perfectly connected" superconcentrators then we can argue that the transformed $\textsc{Transform}(G_N, \text{SC}^{depth}, I, O)$ is $(e/3, (d-e)\cdot d')$-node depth robust — see Theorem 9. In particular, the depth parameter $2d-1$ from Theorem 5 is replaced with $(d-e)\cdot d'$.

We will plug in $d' = \Theta(|E(G_N)|/|V(G_N))$ so that $(d-e)\cdot d' = \Theta((d-e)|E(G_N)|/|V(G_N))$. In section 3.4 we will argue that there is an (e, d)-edge-depth robust graph $e = \Omega(N)$, $d = \Omega(N)$ having $|E(G_N)| = O(N \log N)$ edges. If $|E(G_N)| = O(N \log N)$, $d' = \Omega(\log N)$, and $d = \Omega(N)$ then the final depth parameter $d \cdot d' = \Omega(N \log N) = \Omega(|V(G^*_N)|)$ is linear in the size of the final graph graph G^*_N.

It remains to show (1) how to construct "d'-perfectly connected" superconcentrators and (2) prove the updated depth-robustness lower bound for our transform when we use a family of "d'-perfectly connected" superconcentrators.

We first show how to transform any n-superconcentrator $\text{SC}'(n)$ into a new n-superconcentrator $\text{SC}^{depth}(n, d')$ which also satisfies the property of d-perfect connectivity by adding *at most* $n + d'$ new nodes — see Definition 14. Our transformation is illustrated in Figure 2 with new nodes and edges are depicted in red. More formally $V(\text{SC}^{depth}) = V(\text{SC}) \cup \{z_1, \ldots, z_{d'}\} \cup \{\text{OUT}'_1, \ldots, \text{OUT}'_n\}$ and

$$E(\text{SC}^{depth}) = E(\text{SC}) \cup \{(\text{OUT}'_i, \text{OUT}'_{i+1})\}_{i=1}^{n-1} \cup \{(z_i, z_{i+1})\}_{i=1}^{d'-1} \cup \{(\text{OUT}_1, z_1)\} \cup \{(z_{d'}, \text{OUT}'_1)\} \cup \{(\text{OUT}_i, \text{OUT}'_i)\}_{i=1}^{n}$$

The graph $\text{SC}^{depth}(n, d')$ adds $n + d'$ new nodes and $2n + d' - 1$ new edges to the original superconcentrator $\text{SC}(n)$. It is easy to verify that is both an n-superconcentrator and d'-perfectly connected — see Claim 8 for a formal proof.

The proof that the transformed graph is $(e/3, (d-e)\cdot d')$-node depth robust closely follows our earlier argument in Section 3.2. The crucial difference is that we expand the definitions of internally/externally disrespectful node/edge. Recall that a node $x \in \text{OUT}(\text{SC}_v)$ was internally disrespectful if $depth(x, G^*_N - V S^*) <$

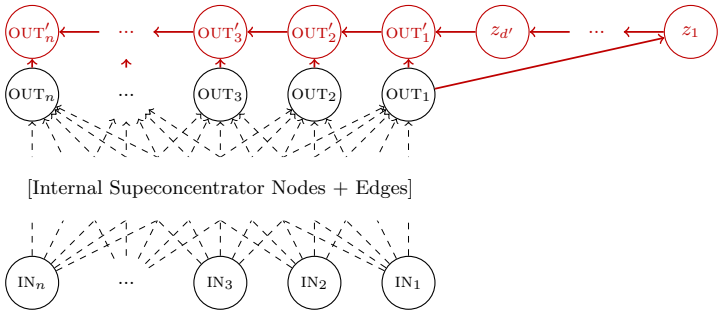


Fig. 2: $\text{SC}^{depth}(n, d')$, a d'-perfectly connected superconcentrator

$r(v, G^*_N, S^*)+1$ because if many input nodes have depth $r(v, G^*_N, S^*)$ we expect that many output nodes will have depth $r(v, G^*_N, S^*)+1$. If $|S^*_v| = 0$ so that SC_v is untouched then we now expect that many output nodes will have depth at least $r(v, G^*_N, S^*) + d'$. Thus, if $|S^*_v| = 0$ and $\texttt{depth}(x, G^*_N - _V S^*) < r(v, G^*_N, S^*) + d'$ then we will still say that the node $x \in \text{OUT}(\text{SC}_v)$ is internally disrespectful and we will say that the corresponding edge $e \in E(G_N)$ such that $\texttt{out}(e) = x$ is also internally disrespectful.

Letting $\text{IntDisrHD}(G^*_N, S^*)$ denote the new set of internally disrespectful nodes we can define $S_{\text{IntDisrHD}}(G^*_N, S^*) \doteq \{e \in E(G_N) \ : \ \texttt{out(e)} \in \text{IntDisrHD}(G^*_N, S^*)\}$ — see definitions 15 and 17. We can still argue that $|\text{IntDisrHD}(G^*_N, S^*) \cap V(\text{SC}_v)| \leq |S^*_v|$ — see Lemma 8.

Similarly, we define a new set, ExtDisrHD, of externally disrespectful nodes in G^*_N for d'-perfectly connected superconcentrators and corresponding set $S_{\text{ExtDisrHD}}$ of edges in G_N. Recall that the prior definition of externally disrespectful nodes captured the intuition that if $(u, v) \in E(G_N)$ then we expect that $r(v, G^*_N, S^*) > r(u, G^*_N, S^*)+1$. However, if $|S^*_u| = 0$ (i.e., SC_u is untouched) then we now expect that $r(v, G^*_N, S^*) \geq r(u, G^*_N, S^*) + d'$. If this is not the case then we will still say that the node $\texttt{out}((u, v)) \in \text{IN}(\text{SC}_v)$ is externally disrespectful along with the edge (u, v) i.e. if either $r(v, G^*_N, S^*) \leq r(u, G^*_N, S^*) + 1$, and $|S^*_u| > 0$ or $r(v, G^*_N, S^*) \leq r(u, G^*_N, S^*) + d'$ and $|S^*_u| = 0$ then the node $x = \texttt{in}((u, v))$ in G^*_N is externally disrespectful. See Definition 16 and Lemma 9 for formal definitions. We can still argue that $|SC_v \cap \text{ExtDisrHD}(G^*_N, S^*)| \leq |S^*_v|$ — see Lemma 9.

This allows us to once again define a set S of $|S| \leq 3|S^*|$ edges to remove from G_N such that $\texttt{depth}(G^*_N - _V S^*) \geq (\texttt{depth}(G_N - _E S) - |S^*|) \cdot d'$. See Definition 17 for the definition of S and prove that $|S| \leq 3|S^*|$ in Lemma 10. Intuitively, if $P = v_1, \ldots, v_d$ is a path in $G_N - _E S$ then we let $U = \{i : |S^*_{v_i}| = 0\}$ denote the set of all untouched nodes on the path. We know that $|U| \geq d - |S^*|$ and for every $i \in U$ we can argue that $r(v_{i+1}, G^*_N, S^*) \geq r(v_i, G^*_N, S^*)+d'$ and we always have $r(v_{i+1}, G^*_N, S^*) \geq r(v_i, G^*_N, S^*)$. Thus, it follows that the $r(v_d, G^*_N, S^*) \geq (d - e) \cdot d'$. See Theorem 9 and its formal proof in Appendix C.

3.4 Edge Depth-Robustness of Grates: Fresh Analysis

The Grates graph [25] has $N' = 2^{n+1} - 1$ nodes and $n2^n$ edges. In this section we provide a new analysis of the Grates graph and demonstrate that the graph is $(e, 2d-1)$-edge depth robust for any parameters e, d such that $e+d \leq 2^n$. Combined with Theorem 9 and an explicit construction of high depth superconcentrators with maximum indegree 2 ,this yields an explicit construction of a (e, d)-node depth robust graph with N nodes, maximum indegree 2 and depth-robustness parameters $e = \Omega(N/\log N)$ and $d = \Omega(N)$.

The Grates graph, first introduced by [25], is a simple, well-studied, edge-depth robust graph. The Grates graph $G_n^{\texttt{Grates}}$ is defined as follows: $V(G_n^{\texttt{Grates}}) \doteq \bigcup_{i=0}^{n}\{0,1\}^i$ and $E(G_n^{\texttt{Grates}}) \doteq \{(u0w, u) \mid |u0w| = n\} \cup \{(u, u1w) \mid |u1w| = n\}$ — see Definition 18. See Figure 3 for an example with $n = 3$ and for an example of least common ancestor e.g., $\mathbf{LCA}(100, 110) = 1$, $\mathbf{LCA}(010, 110) = \epsilon$ and $\mathbf{LCA}(010, 011) = 01$.

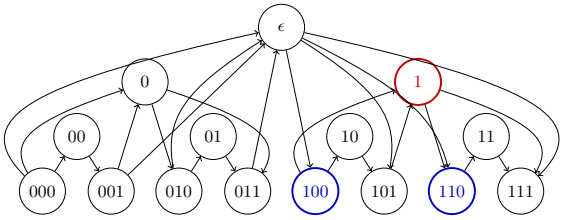


Fig. 3: $G_n^{\texttt{Grates}}$ with $n = 3$

Here it will suffice to make a few observations about $G_n^{\texttt{Grates}}$: (1) $|V(G_n^{\texttt{Grates}})| = 2^{n+1} - 1$ (2) $|E(G_n^{\texttt{Grates}})| = n2^n$ (3) $\texttt{indeg}(G_n^{\texttt{Grates}}) = 2^{n-1}$.

First, we show that Grates is $(e, 2d-1)$-edge-depth robust for any parameters $e+d \leq 2^n$. Intuitively, if we delete $|S| \leq e$ edges from $G_n^{\texttt{Grates}}$ then *at least* $2^n - e$ leaf nodes are untouched. If two leaf nodes x, y are both untouched and x comes before y in a topological ordering of $G_n^{\texttt{Grates}}$ then $G_n^{\texttt{Grates}} -_E S$ will contain the direct path $x \to \mathbf{LCA}(x, y) \to y$. Thus, there is a path containing all $2^n - e$ untouched leaf nodes and *at least* $2^n - e - 1$ intermediate **LCA** nodes. See Appendix D for a formal proof of Theorem 10

Theorem 2. *$G_n^{\texttt{Grates}}$ (defined in Definition 18) is $(e, 2d-1)$-edge depth-robust for all depth-robustness parameters e, d such that $e + d \leq 2^n$.*

Theorem 3 (Informal, See Theorem 10). $\textsc{Transform}(G_n^{\texttt{Grates}}, \mathrm{SC}^{depth}, I, O)$ *is $(e/3, (2n - e - 1)d')$-node depth robust for all $e + d \leq 2^n$, $d > 1$.*

Concrete Parameters: We now briefly examine the concrete depth-robustness parameters that we obtain by applying our transformation to $G_n^{\texttt{Grates}}$. Linear-sized superconcentrators have been shown to exist, including randomized constructions of superconcentrators with $|E(\textsc{sc}(n))|/n)| = 25.3$ [19] and explicit constructions of superconcentrators with $|E(\textsc{sc}(n))|/n = 44$ [1]. We can (explicitly) transform these superconcentrators to obtain new families with maximum

indegree 2 that are d'-perfectly connected obtaining a (non-explicit) family of n-superconcentrators with $27.3n + d'$ nodes and an explicit family with $46 + d'$ nodes.

Given a family of SC^{depth} of d'-perfectly connected superconcentrators with size $|V(\text{SC}^{depth}(N)| \leq cN+d'$ our transformation $\textsc{Transform}(G_n^{\texttt{Grates}}, \text{SC}^{depth}, I, O)$ yields an indegree 2 graph with $|V(G_N^*)| \leq (c+2)n2^n + o(n2^n)$ nodes that is $(\frac{e}{3}, (2 \cdot 2^n - 3e - 1)n)$-*node* depth robust for all $e \leq 2^n$, $d > 1$ — see Theorem 13.

In Lemma 14 and Corollary 1, we instantiate our transformation using the superconcentrator construction of [19]. This yields a graph G_N^* that: (1) has maximum indegree 2, (2) has approximately $29.3n2^n$ nodes, (3) is approximately $(\frac{e}{3}, (2^{n+1} - 3e)n)$-node-depth-robust for any $e \leq 2^n$, and (4) satisfies the cumulative pebbling lower bound $ed \geq \frac{1}{7726.41} \frac{N'^2}{\log N'}$, where $N' = |V(G_N^*)|$. This yields the best currently known concrete lower bound on the ed product, which in turn lower bounds the cumulative pebbling cost [4]. In particular, this improves the corresponding DRSample constant [3] by a factor of approximately 17.7.

3.5 Fractional Depth Robustness

We also show that our transformation preserves *fractional* depth-robustness. Formalized in Appendix B, we give an informal statement of Theorem 6 as follows:

Theorem 4 (Informal, See Theorem 6). *Given a graph G_N which is (e, d, f_1, f_2) fractionally edge depth robust, let $\beta = \frac{1}{3} \cdot \frac{|V(G_N)|}{|V(G_N^*)|} \cdot \min_{v \in V(G_N)} \texttt{mdeg}(v)$. $\textsc{Transform}(G_N, \text{SC}, I, O)$ is $(e/3, 2d-1, f_1 \cdot \beta, f_2 \cdot \beta)$ fractionally node depth robust for all $d > 1$.*

If $\min_{v \in V(G_N)} \texttt{mdeg}(v) = \Theta\Big(\frac{|E(G_N)|}{|V(G_N)|}\Big)$, then $\beta = \Omega(1)$, and the fractional depth-robustness parameters are preserved asymptotically up to constant factors. In particular, this is the case for the Grates construction, where $\min_{v \in V(G_n^{\texttt{Grates}})} \texttt{mdeg}(v) = \Theta(n)$ while $\frac{|E(G_n^{\texttt{Grates}})|}{|V(G_n^{\texttt{Grates}})|} = \Theta(n)$.

If G_N is (e, d, f_1, f_2)-fractionally edge depth robust then, after deleting e edges, at least a f_2 (resp. f_1) fraction of the nodes are endpoints of a length d path. Our main observation is that if node $v \in V(G_N - S)$ is the endpoint of a length d path in G_N then we must $r(v, G_N^*, S^*) \geq 2d - 1$. Thus, most outputs of the superconcentrator SC_v will be endpoints of the path of length $2d-1$. A similar argument applies to lower bound the number of nodes that are the source of a length $2d-1$ path in G_N^* after making a few observations: (1) reversing the edges in a (e, d, f_1, f_2)-edge depth robust graph yields a (e, d, f_2, f_1)-edge depth robust graph i.e., $\texttt{Rev}(G_N)$ is fractionally edge depth robust, (2) reversing the edges in a superconcentrator yields another superconcentrator (Fact 7), and (3) $\texttt{Rev}(\textsc{Transform}(G_N, \text{SC}, I, O)) = \textsc{Transform}(\texttt{Rev}(G_N), \texttt{Rev}(\text{SC}), I', O')$ — see Lemma 6.

Acknowledgments. This work was supported in part by NSF CAREER Award CNS-2047272 and by an Amazon Research Award. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors

and do not necessarily reflect the views of the National Science Foundation or Amazon.

References

1. Alon, N., Capalbo, M.: Smaller explicit superconcentrators. Internet Mathematics **1**(2), 151–163 (2004). https://doi.org/10.1080/15427951.2004.10129083, https://doi.org/10.1080/15427951.2004.10129083
2. Alwen, J., Blocki, J.: Efficiently computing data-independent memory-hard functions. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology – CRYPTO 2016, Part II. Lecture Notes in Computer Science, vol. 9815, pp. 241–271. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2016). https://doi.org/10.1007/978-3-662-53008-5_9
3. Alwen, J., Blocki, J., Harsha, B.: Practical graphs for optimal side-channel resistant memory-hard functions. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017: 24th Conference on Computer and Communications Security. pp. 1001–1017. ACM Press, Dallas, TX, USA (Oct 31 – Nov 2, 2017). https://doi.org/10.1145/3133956.3134031
4. Alwen, J., Blocki, J., Pietrzak, K.: Depth-robust graphs and their cumulative memory complexity. In: Coron, J.S., Nielsen, J.B. (eds.) Advances in Cryptology – EUROCRYPT 2017, Part III. Lecture Notes in Computer Science, vol. 10212, pp. 3–32. Springer, Cham, Switzerland, Paris, France (Apr 30 – May 4, 2017). https://doi.org/10.1007/978-3-319-56617-7_1
5. Alwen, J., Blocki, J., Pietrzak, K.: Sustained space complexity. In: Nielsen, J.B., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2018, Part II. Lecture Notes in Computer Science, vol. 10821, pp. 99–130. Springer, Cham, Switzerland, Tel Aviv, Israel (Apr 29 – May 3, 2018). https://doi.org/10.1007/978-3-319-78375-8_4
6. Alwen, J., Serbinenko, V.: High parallel complexity graphs and memory-hard functions. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th Annual ACM Symposium on Theory of Computing. pp. 595–603. ACM Press, Portland, OR, USA (Jun 14–17, 2015). https://doi.org/10.1145/2746539.2746622
7. Biryukov, A., Dinu, D., Khovratovich, D.: Argon2: New generation of memory-hard functions for password hashing and other applications. In: 2016 IEEE European Symposium on Security and Privacy. pp. 292–302. IEEE Computer Society Press, Saarbrücken, Germany (Mar 21–24, 2016). https://doi.org/10.1109/EuroSP.2016.31
8. Blocki, J., Cinkoske, M.: A new connection between node and edge depth robust graphs. In: Lee, J.R. (ed.) ITCS 2021: 12th Innovations in Theoretical Computer Science Conference. vol. 185, pp. 64:1–64:18. Leibniz International Proceedings in Informatics (LIPIcs), Virtual Conference (Jan 6–8, 2021). https://doi.org/10.4230/LIPIcs.ITCS.2021.64
9. Blocki, J., Cinkoske, M., Lee, S., Son, J.Y.: On explicit constructions of extremely depth robust graphs (2022), https://arxiv.org/abs/2110.04190
10. Blocki, J., Harsha, B., Kang, S., Lee, S., Xing, L., Zhou, S.: Data-independent memory hard functions: New attacks and stronger constructions. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology – CRYPTO 2019, Part II. Lecture Notes in Computer Science, vol. 11693, pp. 573–607. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 18–22, 2019). https://doi.org/10.1007/978-3-030-26951-7_20

11. Blocki, J., Holman, B.: Sustained space and cumulative complexity trade-offs for data-dependent memory-hard functions. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022, Part III. Lecture Notes in Computer Science, vol. 13509, pp. 222–251. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15982-4_8
12. Blocki, J., Holman, B.: Towards practical data-dependent memory-hard functions with optimal sustained space trade-offs in the parallel random oracle model. In: CRYPTO. Springer (2026)
13. Blocki, J., Lee, S., Zhou, S.: Approximating cumulative pebbling cost is unique games hard. In: Vidick, T. (ed.) ITCS 2020: 11th Innovations in Theoretical Computer Science Conference. vol. 151, pp. 13:1–13:27. Leibniz International Proceedings in Informatics (LIPIcs), Seattle, WA, USA (Jan 12–14, 2020). https://doi.org/10.4230/LIPIcs.ITCS.2020.13
14. Blocki, J., Zhou, S.: On the depth-robustness and cumulative pebbling cost of Argon2i. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017: 15th Theory of Cryptography Conference, Part I. Lecture Notes in Computer Science, vol. 10677, pp. 445–465. Springer, Cham, Switzerland, Baltimore, MD, USA (Nov 12–15, 2017). https://doi.org/10.1007/978-3-319-70500-2_15
15. Blocki, J., Zhou, S.: On the computational complexity of minimal cumulative cost graph pebbling. In: Meiklejohn, S., Sako, K. (eds.) FC 2018: 22nd International Conference on Financial Cryptography and Data Security. Lecture Notes in Computer Science, vol. 10957, pp. 329–346. Springer Berlin Heidelberg, Germany, Nieuwpoort, Curaçao (Feb 26 – Mar 2, 2018). https://doi.org/10.1007/978-3-662-58387-6_18
16. Dziembowski, S., Faust, S., Kolmogorov, V., Pietrzak, K.: Proofs of space. In: Gennaro, R., Robshaw, M.J.B. (eds.) Advances in Cryptology – CRYPTO 2015, Part II. Lecture Notes in Computer Science, vol. 9216, pp. 585–605. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 2015). https://doi.org/10.1007/978-3-662-48000-7_29
17. Erdös, P., Graham, R., Szemerédi, E.: On sparse graphs with dense long paths. Computers & Mathematics with Applications **1**(3), 365–369 (1975). https://doi.org/https://doi.org/10.1016/0898-1221(75)90037-1, https://www.sciencedirect.com/science/article/pii/0898122175900371
18. Gabber, O., Galil, Z.: Explicit constructions of linear size superconcentrators. In: 20th Annual Symposium on Foundations of Computer Science (FOCS 1979). pp. 364–370. IEEE Computer Society (1979)
19. Kolmogorov, V., Rolínek, M.: Superconcentrators of density 25.3. CoRR **abs/1405.7828** (2014), http://arxiv.org/abs/1405.7828
20. Lev, G., Valiant, L.: Size bounds for superconcentrators. Theoretical Computer Science **22**(3), 233–251 (1983). https://doi.org/https://doi.org/10.1016/0304-3975(83)90105-6, https://www.sciencedirect.com/science/article/pii/0304397583901056
21. Mahmoody, M., Moran, T., Vadhan, S.P.: Publicly verifiable proofs of sequential work. In: Kleinberg, R.D. (ed.) ITCS 2013: 4th Innovations in Theoretical Computer Science. pp. 373–388. Association for Computing Machinery, Berkeley, CA, USA (Jan 9–12, 2013). https://doi.org/10.1145/2422436.2422479
22. Pietrzak, K.: Proofs of catalytic space. In: Blum, A. (ed.) ITCS 2019: 10th Innovations in Theoretical Computer Science Conference. vol. 124, pp. 59:1–59:25. Leibniz International Proceedings in Informatics (LIPIcs), San Diego, CA, USA (Jan 10–12, 2019). https://doi.org/10.4230/LIPIcs.ITCS.2019.59

23. Pippenger, N.: Superconcentrators. SIAM Journal on Computing **6**(2), 298–304 (1977). https://doi.org/10.1137/0206022, https://doi.org/10.1137/0206022
24. Schnitger, G.: A family of graphs with expensive depth-reduction. Theoretical Computer Science **18**(1), 89–93 (1982). https://doi.org/https://doi.org/10.1016/0304-3975(82)90113-X, https://www.sciencedirect.com/science/article/pii/030439758290113X
25. Schnitger, G.: On depth-reduction and grates. In: Proceedings of the 24th Annual Symposium on Foundations of Computer Science. p. 323–328. SFCS '83, IEEE Computer Society, USA (1983). https://doi.org/10.1109/SFCS.1983.38, https://doi.org/10.1109/SFCS.1983.38
26. Schöning, U.: Smaller superconcentrators of density 28. Information Processing Letters **98**(4), 127–129 (2006). https://doi.org/https://doi.org/10.1016/j.ipl.2006.01.006, https://www.sciencedirect.com/science/article/pii/S0020019006000159
27. Valiant, L.G.: On non-linear lower bounds in computational complexity. In: Proceedings of the Seventh Annual ACM Symposium on Theory of Computing. p. 45–53. STOC '75, Association for Computing Machinery, New York, NY, USA (1975). https://doi.org/10.1145/800116.803752, https://doi.org/10.1145/800116.803752

A Proofs of Section 3.1

A.1 Reminder of Our Transformation

Given any (edge-depth-robust) graph with N nodes G_N, we define a graph G_N^* as follows: For each vertex $v \in V(G_n)$ of the graph, we select an superconcentrator and replace v with the superconcentrator. We call this superconcentrator SC_v. We then connect every edge between two nodes to come from an output of the source node's superconcentrator and lead to an input of the destination node's superconcentrator.

Definition 3. *Given a graph $G_n = (V, E)$, and a family of superconcentrators $\text{SC} = \{\text{SC}(i)\}_{i=1}^{\texttt{mdeg}(G_n)}$, for all vertices $v \in V(G_n)$, let $\text{SC}_v \doteq \text{SC}(\texttt{mdeg}(G_n, v))$ denote a fresh copy of $\text{SC}((\texttt{mdeg}(G_n, v))$. We let $u_1, \ldots, u_{\texttt{indeg}(v)}$ (resp. $w_1, \ldots, w_{\texttt{outdeg}(v)}$) denote the parents (resp. children) of node $v \in V(G_n)$ and let $\texttt{in}_v : \{u_1, \ldots, u_{\texttt{indeg}(v)}\} \to \text{IN}(\text{SC}_v)$ (resp. $\texttt{out}_v : \{w_1, \ldots, w_{\texttt{outdeg}(v)}\} \to \text{OUT}(\text{SC}_v)$) denote a one-to-one function.*

Fixing families $I = \{\texttt{in}_v\}_{v \in V}$ and $O = \{\texttt{out}_v\}_{v \in V}$ of such one-to-one functions we define the graph $\textsc{Transform}(G_N, \text{SC}, I, O) = G_N^* = (V_n^*, E_n^*)$ as follows:

$$V_n^* \doteq \left(\bigcup_{v \in V(G_n)} V(\text{SC}_v) \right) \quad E_n^* \doteq \left(\bigcup_{v \in V(G_n)} E(\text{SC}_v) \right) \cup \left(\bigcup_{e \in E(G_n)} (\texttt{out}(e), \texttt{in}(e)) \right)$$

where we define $\texttt{in} : E(G_n) \to \bigcup_{v \in V(G_n)} \text{IN}(\text{SC}_v)$ as $\texttt{in}((u, v)) = \texttt{in}_v(u)$ and $\texttt{out} : E(G_n) \to \bigcup_{v \in V(G_n)} \text{OUT}(\text{SC}_v)$ as $\texttt{out}((v, w)) = \texttt{out}_v(w)$.

When G_N, SC, I, and O are arbitrary or clear from context, we often simply write G_N^* for brevity. Notably, while the size of the resulting graph G_N^* depends

upon n and the superconcentrators used, if the superconcentrator constructions used have constant indegree, so too will G^*_N.

A.2 Proving Our Transformation Preserves Depth Robustness

We now wish to show that the resulting graph $G^*_N = \textsc{Transform}(G_N, \mathrm{SC}, I, O)$ is node-Depth Robust as long as the original graph is edge-Depth Robust. To do so, we will consider removing any set of nodes S^* from G^*_N. We will show a mapping from $S^* \in V(G^*_N)$ to a set of edges $S \in E(G_N)$ of size at most $|S| \leq 3|S^*|$. Assuming that $|S^*| \leq e/3$ and G_N is (e,d)-edge depth robust it follows that there is some node $u \in V_N$ with depth $\mathtt{depth}(u, G^*_N - {}_E S) \geq d$. We will argue that the corresponding output nodes $x \in \textsc{out}(\mathrm{SC}_u)$ have depth at least $2 \cdot \mathtt{depth}(u, G_N - {}_E S) \geq 2d - 1$. It immediately follows that G^*_N is $(e/3, 2d-1)$-node depth robust. We say that for any removed node $x \in S^*$, $\mathtt{depth}(x, G^*_N - {}_V S^*) = 0$

Reminder of Definition 2. *Given any superconcentrator* SC_v *in* G^*_N *and a set* S^* *of nodes removed from* G^*_N*. We define the set* $S^*_v \doteq S^* \cap \mathrm{SC}_v$ *as the set of all nodes removed from that superconcentrator and we define* $\gamma_v(S^*) = |S^*_v|/|\textsc{in}(\mathrm{SC}_v)|$*. When* S^* *is clear from context we will simply write* γ_v*.*

Definition 4. *Given a superconcentrator* SC_v *in* G^*_N*, a set* S^* *of nodes removed from* G^*_N*, and a threshold fraction* γ *we define the function* $r(v, G^*_N, \gamma, S^*) = \max\{t \ : \ |\{x \in \textsc{in}(\mathrm{SC}_v) \ : \ \mathtt{depth}(x, G^*_N - {}_V S^*) \geq t\}| > \gamma\,|\textsc{in}(\mathrm{SC}_v)|\}$ *to find the largest depth such that at least* γ *fraction of the inputs to the superconcentrator have that depth or greater. If* $\gamma \geq 1$*, we define* $r(v, G^*_N, \gamma, S^*)$ *to be* 0*. We will also define* $r(v, G^*_N, S^*) = r(v, G^*_N, \gamma_v(S^*), S^*)$ *since we will always set* $\gamma = \gamma_v(S^*)$ *in our analysis.*

Definition 5. *Given a graph* G_N*, a family of superconcentrators* SC*, and families* $I = \{\mathtt{in}_v\}_{v \in V}$ *and* $O = \{\mathtt{out}_v\}_{v \in V}$ *of one-to-one functions, let* $G^*_N = \textsc{Transform}(G_N, \mathrm{SC}, I, O)$ *be the transformed graph constructed as in Definition 3. Given a set* S^* *of nodes removed from* G^*_N*, and a superconcentrator* SC_v *in* G^*_N*, we define the predicate* $\textsc{Spoiled}(v, G^*_N, S^*)$ *to be* 1 *if and only if* $\gamma_v(S^*) \geq \frac{2}{3}$ *— i.e., the number of nodes removed from* SC_v *is at least two-thirds the size of the input to* SC_v *— and* 0*, otherwise.*

We say that a superconcentrator is spoiled if $\textsc{Spoiled}(v, G^*_N, S^*)$ is true. A spoiled superconcentrator is no longer useful since it is possible that the entire input or output set (or some other combination, potentially including internal nodes) was removed, and there may no longer be any useful paths through the superconcentrator. Instead, we will exclude such superconcentrators, removing all incoming and outgoing edges so that it is not used in any meaningful paths.

Definition 6. *Given a graph* G_N*, a family of superconcentrators* SC*, and families* $I = \{\mathtt{in}_v\}_{v \in V}$ *and* $O = \{\mathtt{out}_v\}_{v \in V}$ *of one-to-one functions, let* $G^*_N = \textsc{Transform}(G_N, \mathrm{SC}, I, O)$ *be the transformed graph constructed as in Definition 3. Given a set* S^* *of nodes removed from* G^*_N*, we define the set* $\textsc{ToExclude}(G_N, \mathrm{SC}, I, O, S^*)$

of edges in G_N (the original graph) as the set of all edges into or out of spoiled superconcentrators, i.e. $\textsc{ToExclude}(G_N, \mathrm{SC}, I, O, S^*) = \{(u,v) \in E(G_N) : \textsc{Spoiled}(u, G_N^*, S^*) = 1 \vee \textsc{Spoiled}(v, G_N^*, S^*) = 1\}$

Since such edges are removed from the original graph, we must bound the size $|\textsc{ToExclude}(G_N, \mathrm{SC}, I, O, S^*)|$ in relation to the size $|S^*|$ in order to show the relationship between the edge-depth robustness of G_N and the node-depth robustness of G_N^*.

Lemma 1. *Given a graph G_N, a family of superconcentrators* SC*, and families $I = \{\mathtt{in}_v\}_{v\in V}$ and $O = \{\mathtt{out}_v\}_{v\in V}$ of one-to-one functions, let $G_N^* = \textsc{Transform}(G_N, \mathrm{SC}, I, O)$ be the transformed graph constructed as in Definition 3. Given a set S^* of nodes removed from G_N^*, $|\textsc{ToExclude}(G_N, \mathrm{SC}, I, O, S^*)| \leq 3|S^*|$*

Proof. For all $v \in V(G_N)$, $|\textsc{in}(\textsc{sc}_v)| = |\textsc{out}(\textsc{sc}_v)|$, so if $\textsc{sc}_v$ is spoiled at most $2|\textsc{in}(\textsc{sc}_v)|$ additional edges $E_v^+ = \{(a,b) \in E(G_N) : a = v \vee b = v$ are added to $\textsc{ToExclude}(G_N, \mathrm{SC}, I, O, S^*)$. On the other hand, if $\textsc{sc}_v$ is spoiled then $\gamma_v = |S_v^*|/|\textsc{in}(\textsc{sc}_v)| \geq \frac{2}{3}$. Thus, for all $v \in V(G_N)$, $|E_v^+| \leq 2\frac{3}{2}|S_v^*| = 3|S_v^*|$. Therefore, $|\textsc{ToExclude}(G_N, \mathrm{SC}, I, O, S^*)| \leq \bigcup_{v\in V(G_n)} |E_v^+| \leq \bigcup_{v\in V(G_n)} 3|S_v^*| \leq 3|S^*|$. □

Definition 7. *Given a graph G_N, a family of superconcentrators* SC*, and families $I = \{\mathtt{in}_v\}_{v\in V}$ and $O = \{\mathtt{out}_v\}_{v\in V}$ of one-to-one functions, let $G_N^* = \textsc{Transform}(G_N, \mathrm{SC}, I, O)$ be the transformed graph constructed as in Definition 3. Given a superconcentrator $\textsc{sc}_v$ in G_N^* and a set S^* of nodes removed from G_N^*, we define the set $\textsc{Dest}(v, G_N^*, S^*) = \textsc{in}(\textsc{sc}_v) \cap S_v^*$ to be the set of all input nodes which are removed. We further define the set $\textsc{Dest}(G_N^*, S^*) = \bigcup_{v\in V(G_N)} \textsc{Dest}(v, G_N^*, S^*)$ to be all input nodes removed in the graph G_N^*.*

We say that an input node is destroyed if it is in the set $\textsc{Dest}_{\textsc{in}}(v, G_N^*, S^*)$, and that it is present, otherwise. It follows immediately from the definition that $|\textsc{Dest}_{\textsc{in}}(v, G_N^*, S^*)| \leq |S_v^*|$.

Since every output node in a superconcentrator has a path from the highest-depth input node(s) in a superconcentrator, one might expect that in any superconcentrator $\textsc{sc}_v$, for every output node $u \in \textsc{out}(\textsc{sc}_v)$, the depth will increase, i.e. $\mathtt{depth}\,(u, G_N^* - _V S^*) \geq r(v, G_N^*, S^*)+1$. However, due to the removal of nodes in S^* this may no longer be the case, and some output vertices may "disrespect" the expected increasing order of depths. We call such output nodes "internally disrespectful" and formalize this notion as follows:

Definition 8. *Given a graph G_N, a family of superconcentrators* SC*, and families $I = \{\mathtt{in}_v\}_{v\in V}$ and $O = \{\mathtt{out}_v\}_{v\in V}$ of one-to-one functions, let $G_N^* = \textsc{Transform}(G_N, \mathrm{SC}, I, O)$ be the transformed graph constructed as in Definition 3. Given a superconcentrator $\textsc{sc}_v$ in G_N^* and a set S^* of nodes removed from G_N^*, we define the set $\textsc{IntDisr}(v, G_N^*, S^*)$ to be the set of all output nodes which do not have depth at least $r(v, G_N^*, S^*) + 1$, i.e. $\textsc{IntDisr}(v, G_N^*, S^*) \doteq \{x \in \textsc{out}(\textsc{sc}_v) : \mathtt{depth}(x, G_N^* - _V S^*) < r(v, G_N^*, S^*) + 1\}$. We further define the set $\textsc{IntDisr}(G_N^*, S^*) = \bigcup_{v\in V(G_N)} \textsc{IntDisr}(v, G_N^*, S^*)$ of all internally disrespectful nodes in the graph G_N^*.*

We also upper bound the number of internally disrespectful nodes in a given superconcentrator.

Lemma 2. *Given a graph G_N, a family of superconcentrators* SC*, and families $I = \{\mathtt{in}_v\}_{v\in V}$ and $O = \{\mathtt{out}_v\}_{v\in V}$ of one-to-one functions, let $G_N^* = \textsc{Transform}(G_N, \text{SC}, I, O)$ be the transformed graph constructed as in Definition 3. Given a superconcentrator SC_v in G_N^* and a set S^* of nodes removed from G_N^*, $|\textsc{IntDisr}(v, G_N^*, S^*)| \leq |S_v^*|$*

Proof. Suppose, for the sake of contradiction, that there are more than $|S_v^*|$ nodes in $\textsc{IntDisr}(v, G_N^*, S^*)$. Then there exists a set $B \subseteq \textsc{IntDisr}(v, G_N^*, S^*)$ of at least $|B| \geq |S_v^*|+1$ internally disrespectful output vertices. By definition 4 there exists a set $\textsc{Good}_{\text{IN}}$ containing at least $\gamma_v|\text{IN}(\text{SC}_v)| + 1 = |S_v^*| + 1$ input nodes in SC_v such that $\forall x \in \textsc{Good}_{\text{IN}}, \mathtt{depth}(x, G_N^* -_V S^*) \geq r(v, G_N^*, S^*)$. Because SC_v is a superconcentrator there are at least $|S_v^*|+1$ node disjoint paths connecting $\textsc{Good}_{\text{IN}}$ to B. Yet there are at most $|S_v^*|$ vertices removed from these paths. Thus, there must remain at least 1 path from some node $a \in \textsc{Good}_{\text{IN}}$ to some node $b \in B$ which does not contain any vertex in S_v^*. Yet this would mean the vertex b has depth at least $r(v, G_N^*, S^*) + 1$, since a has depth at least $r(v, G_N^*, S^*)$ by definition of the set $\textsc{Good}_{\text{IN}}$. However, if b has depth at least $r(v, G_N^*, S^*) + 1$ then by definition of $\textsc{IntDisr}(v, G_N^*, S^*)$ we would have $b \notin \textsc{IntDisr}(v, G_N^*, S^*)$. Contradiction! It follows that there can be at most $|S_v|$ nodes in $\textsc{IntDisr}(v, G_N^*, S^*)$. □

There is one other way in which the expected order of increasing r values may be disrupted. Suppose that there is an edge $e = (u, v)$ in the original graph G_n which corresponds to an edge $e^* = (a, b)$ in G_n^* between some node $a \in \text{OUT}(\text{SC}_u)$ and some node $b \in \text{IN}(\text{SC}_v)$. Because the edge (u, v) exists in the original graph we typically expect that $r(v, G_n^*, S^*) \geq 2 + r(u, G_n^*, S^*)$. If this expected depth increase property does not hold (and a is not internally disrespectful and b has not been deleted) we say that the node b is "externally disrespectful" we formalize this as follows.

Definition 9. *Given a graph G_N, a family of superconcentrators* SC*, and families $I = \{\mathtt{in}_v\}_{v\in V}$ and $O = \{\mathtt{out}_v\}_{v\in V}$ of one-to-one functions, let $G_N^* = \textsc{Transform}(G_N, \text{SC}, I, O)$ be the transformed graph constructed as in Definition 3. Given a superconcentrator SC_v in G_N^* and a set S^* of nodes removed from G_N^*, for any node $b \in \text{IN}(\text{SC}_v)$, we say that b is "externally disrespectful" if $\exists u \in V(G_N)$ such that $\mathtt{In}_v(u) = b$ (where $\mathtt{In}_v$ is the function defined in Definition 3) and (1) $b \notin \textsc{Dest}(v, G_N^*, S^*)$, (2) $r(v, G_N^*, S^*) \leq r(u, G_N^*, S^*)+2$, (3) The node $a = \mathtt{out}(u, v)$ is not internally disrespectful, i.e. $a \notin \textsc{IntDisr}(u, G_N^*, S^*)$, and (4) neither v nor u is spoiled, i.e. $\textsc{Spoiled}(v, G_N^*, S^*) = 0 \wedge \textsc{Spoiled}(u, G_N^*, S^*) = 0$.*

We define the set $\textsc{ExtDisr}(v, G_N^*, S^*)$ to be the set of all externally disrespectful nodes in $\text{IN}(\text{SC}_v)$. We further define the set $\textsc{ExtDisr}(G_N^*, S^*) = \bigcup_{v\in V(G_N)} \textsc{ExtDisr}(v, G_N^*, S^*)$ of all externally disrespectful nodes in the graph G_N^*.

We upper bound the number of externally disrespectful nodes in a given superconcentrator as follows.

Lemma 3. *Given a superconcentrator* SC_v *in* G_N^* *and a set* S^* *of nodes removed from* G_N^*, $|\textsc{ExtDisr}(v, G_N^*, S^*)| \leq |S_v^*|$.

Proof. Suppose, for the sake of contradiction, that $|\textsc{ExtDisr}(v, G_N^*, S^*)| > |S_v^*|$. Let $\rho = r(v, G_N^*, S^*)$ be the r value of the given superconcentrator SC_v. For all $x \in \textsc{ExtDisr}(v, G_N^*, S^*)$, let $u_x \in V_n$ be the unique node such that $\texttt{in}_v(u_x) = x$ and let $a_x = \mathit{out}(u_x, v)$. Observe that if such a u_x does not exist, then x cannot be externally disrespectful, by definition. By property (2) of an externally disrespectful node, we know that $r(u_x, G_N^*, S^*) \geq \rho - 1$. We know by property (3) of an externally disrespectful node that none of the nodes a_x are internally disrespectful. Observe that this also means that a_x cannot be in the set S^* of removed nodes since any output node in S^* will be internally disrespectful by definition. Furthermore, by property (1) of an externally disrespectful node the node $x \notin S^*$ i.e., x was not deleted.

This means that $\forall x \in \textsc{ExtDisr}(v, G_N^*, S^*)$ we have

$$\begin{aligned} \texttt{depth}(x, G_N^* - _V S^*) &\geq \texttt{depth}(a_x, G_N^* - _V S^*) + 1 && (1)\\ &\geq r(u_x, G_N^*, S^*) + 2 && (2)\\ &\geq \rho + 1\ . && (3) \end{aligned}$$

Where inequality 1 follows because both nodes $x, a_X \notin S^*$ have not been deleted. Inequality 2 follows because a_x is not internally disrespectful, and inequality 3 from property (2) of an externally disrespectful node. Thus, there are at least $|S_v^*| + 1$ nodes x with $\texttt{depth}(x, G_N^* - _V S^*) \geq \rho + 1$. It follows from Definition 4 that $r(v, G_n^*, S^*) \geq \rho + 1$. Since we defined $\rho = r(v, G_n^*, S^*)$ this implies that $r(v, G_n^*, S^*) \geq 1 + r(v, G_n^*, S^*)$ a clear contradiction. This contradiction means our original assumption is impossible, and $|\textsc{ExtDisr}(v, G_N^*, S^*)| \leq |S_v^*|$. □

Now that we have defined and bounded these four sets of bad nodes for each superconcentrator, $\textsc{Spoiled}(v, G_N^*, S^*)$ $\textsc{Dest}(v, G_N^*, S^*)$, $\textsc{IntDisr}(v, G_N^*, S^*)$, and $\textsc{ExtDisr}(v, G_N^*, S^*)$, we will use these sets to define the set S of edges which must be removed from the original graph G_N.

Definition 10. *Given* G_N *let* $G_N^* = \textsc{Transform}(G_N, \text{SC}, I, O)$ *be the transformed graph constructed as in Definition 3 and let* $S^* \in V(G_N^*)$ *be any set of nodes. In addition to the set* $\textsc{ToExclude}(G_N^*, S^*) \subseteq E(G_N)$ *we define three additional sets of edges to remove from* G_N *(1)* $S_{\textsc{Dest}} \doteq \{e \in E(G_N) \ : \ \texttt{in}(\texttt{e}) \in \textsc{Dest}(G_N^*, S^*)\}$, *(2)* $S_{\textsc{IntDisr}} \doteq \{e \in E(G_N) \ : \ \texttt{out}(\texttt{e}) \in \textsc{IntDisr}(G_N^*, S^*)\}$, *and (3)* $S_{\textsc{ExtDisr}} \doteq \{e \in E(G_N) \ : \ \texttt{in}(\texttt{e}) \in \textsc{ExtDisr}(G_N^*, S^*)\}$. *We further define the set of all removed edges* $S = \textsc{NtoE}(G_N^*, S^*) \doteq \textsc{ToExclude}(G_N^*, S^*) \cup S_{\textsc{Dest}} \cup S_{\textsc{IntDisr}} \cup S_{\textsc{ExtDisr}}$.

We can bound the total size of this set S.

Lemma 4. *For any graph G^*_N and any subset $S^* \subseteq V(G^*_N)$ of nodes we have* $|\textsc{NtoE}(G^*_N, S^*)| \leq 3|S^*|$.

Proof. Let $\texttt{Spoiled} \subset V(G_N)$ denote the set of all nodes $v \in V(G_N)$ for which $\textsc{Spoiled}(v, G^*_N, S^*) = 1$ i.e., the corresponding super concentrator $\textsc{sc}_v$ in G^*_N was 'spoiled' by removing at least $|\textsc{in}(\textsc{sc}_v)|$ nodes from $\textsc{sc}_v$. We now partition the set S^* into two disjoint subsets, $S^*_A \doteq \bigcup_{v \in \texttt{Spoiled}} S^*_v$ and $S^*_B \doteq \bigcup_{v \notin \texttt{Spoiled}} S^*_v = S^* \setminus S^*_A$ — recall that $S^*_v = \textsc{sc}_v \cap S^*$. Let $S = \textsc{NtoE}(G^*_N, S^*)$ be the corresponding set of edges. We can similarly partition S into two sets $S_A = \{(u,v) \in S \ : \ u \in \texttt{Spoiled} \vee v \in \texttt{Spoiled}\}$ and $S_B = \{(u,v) \in S \ : \ u \notin \texttt{Spoiled} \wedge v \notin \texttt{Spoiled}\}$.

We will show that $|S_A| \leq 3\,|S^*_A|$ and $|S_B| \leq 3\,|S^*_B|$. From this it immediately follows that

$$|S| = |S_A| + |S_B| \leq 3\,|S^*_A| + 3\,|S^*_B| \leq 3|S^*| \ .$$

We first note that $S_A = \textsc{ToExclude}(G_N, \textsc{sc}, I, O, S^*_A)$ so it follows immediately from Lemma 1 that $|S_A| \leq 3|S^*_A|$. We now bound $|S_B|$. We first observe that $S_B = (S_{\textsc{Dest}} \cup S_{\textsc{IntDisr}} \cup S_{\textsc{ExtDisr}}) \setminus \textsc{ToExclude}(G_N, \textsc{sc}, I, O, S^*_A)$. Thus, it suffices to show that $|S_{\textsc{Dest}} \setminus S_A| \leq |S^*_B|$ and $|S_{\textsc{IntDisr}} \setminus S_A| \leq |S^*_B|$ and $|S_{\textsc{IntDisr}} \setminus S_A| \leq |S^*_B|$.

Given $v \in V(G_N)$ let $S_{\textsc{Dest},v} \doteq \{(x,y) \in S_{\textsc{Dest}} \setminus S_A : y \in \textsc{in}(\textsc{sc}_v)\}$ be the set of edges in $S_{\textsc{Dest}}$ which were destroyed because we removed an input node from $\textsc{sc}_v$. For every node $v \in V(G_N)$ we have $|S_{\textsc{Dest},v}| \leq |S^*_v| = \gamma_v\,|\textsc{in}(\textsc{sc}_v)|$. Because any edge $(x,y) \in S_{\textsc{Dest}} \setminus S_A$ must have both x and y unspoiled we also have

$$S_{\textsc{Dest}} \setminus S_A = \bigcup_{v \in \overline{\texttt{Spoiled}}} S_{\textsc{Dest},v} \ .$$

It follows immediately that

$$|S_{\textsc{Dest}} \setminus S_A| \leq \sum_{v \in \overline{\texttt{Spoiled}}} |S_{\textsc{Dest},v}| \leq \sum_{v \in \overline{\texttt{Spoiled}}} \gamma_v\,|\textsc{in}(\textsc{sc}_v)| = \sum_{v \in \overline{\texttt{Spoiled}}} |S^*_v| = |S^* \setminus S_A| = |S^*_B|$$

Similarly, given $v \in V(G_N)$ let $S_{\textsc{IntDisr},v} \doteq \{(x,y) \in S_{\textsc{IntDisr}} \setminus S_A : x \in \textsc{out}(\textsc{sc}_v)\}$ be the set of edges in $S_{\textsc{IntDisr}}$ which are internally disrespectful and exit the superconcentrator $\textsc{sc}_v$. For every node $v \in V(G_N)$ we have

$$|S_{\textsc{IntDisr},v}| \leq |\textsc{IntDisr}(v, G^*_N, S^*)| \leq |S^*_v| \leq \gamma_v\,|\textsc{in}(\textsc{sc}_v)| \ .$$

The first inequality follows because each node in $\textsc{IntDisr}(v, G^*_N, S^*) \subseteq \textsc{out}(\textsc{sc}_v)$ is incident to exactly one edge from $S_{\textsc{IntDisr},v}$, the second inequality follows by Lemma 2 and the third follows by definition. We now have

$$S_{\textsc{IntDisr}} \setminus S_A = \bigcup_{v \in \overline{\texttt{Spoiled}}} S_{\textsc{IntDisr},v}$$

because any edge $(x,y) \in S_{\textsc{IntDisr}} \setminus S_A$ must have both x and y unspoiled. Thus, it immediately follows that

$$|S_{\textsc{IntDisr}} \setminus S_A| \leq \sum_{v \in \overline{\texttt{Spoiled}}} |S_{\textsc{IntDisr},v}| \leq \sum_{v \in \overline{\texttt{Spoiled}}} \gamma_v\,|\textsc{in}(\textsc{sc}_v)| = \sum_{v \in \overline{\texttt{Spoiled}}} |S^*_v| = |S^* \setminus S_A| = |S^*_B| \ .$$

Finally, given $v \in V(G_N)$ let $S_{\text{ExtDisr},v} \doteq \{(x,y) \in S_{\text{ExtDisr}} \setminus S_A : y \in \text{IN}(\text{SC}_v)\}$ be the set of edges in S_{ExtDisr} which are externally disrespectful and enter the superconcentrator SC_v. For every node $v \in V(G_N)$ we have

$$|S_{\text{ExtDisr},v}| \leq |\text{ExtDisr}(v, G_N^*, S^*)| \leq |S_v^*| \leq \gamma_v \, |\text{IN}(\text{SC}_v)| \ .$$

The first line follows because each node in $\text{ExtDisr}(v, G_N^*, S^*) \subseteq \text{IN}(\text{SC}_v)$ is incident to exactly one edge from S_{ExtDisr}, the second inequality follows directly from Lemma 3 and the third follows by definition. We now have

$$S_{\text{ExtDisr}} \setminus S_A = \bigcup_{v \in \overline{\texttt{Spoiled}}} S_{\text{ExtDisr},v}$$

because any edge $(x,y) \in S_{\text{ExtDisr}} \setminus S_A$ must have both x and y unspoiled. Thus, it immediately follows that

$$|S_{\text{ExtDisr}} \setminus S_A| \leq \sum_{v \in \overline{\texttt{Spoiled}}} |S_{\text{ExtDisr},v}| \leq \sum_{v \in \overline{\texttt{Spoiled}}} \gamma_v \, |\text{IN}(\text{SC}_v)| = \sum_{v \in \overline{\texttt{Spoiled}}} |S_v^*| = |S^* \setminus S_A| = |S_B^*| \, .$$

□

It now remains to show the actual connection between this set S and S^*, intuitively that removing S^* from G_N^* is (nearly) equivalent to removing S from G_N. First, we give a helpful lemma which contains the crux of our argument:

Lemma 5. *Given a graph G_N, a family of superconcentrators* SC *and families of one-to-one functions I and O as specified in definition 3, let $G_N^* = \text{Transform}(G_N, \text{SC}, I, O)$ and let $S^* \subseteq V(G_N^*)$ be a subset of nodes. Suppose that $G_N - \text{NtoE}(G_N^*, S^*)$ has a path of length d ending at some node $v_d \in V(G_N)$. Then $G_N^* - S^*$ contains a path of length $2d-2$ ending at a node $u'_d \in \text{IN}(\text{SC}_{v_v})$ with $r(v_d, G_N^*, S^*) \geq 2d-2$.*

Proof. Let $P(v_d) = v_1, \ldots, v_d$ be the path of length d in $G_N - {}_E\text{NtoE}(G_N^*, S^*)$. By definition, we have $(v_i, v_{i+1}) \in E(G_N - {}_E S)$ for all $i \in [1, d-1]$. Define $r_i \doteq r(v_i, G_N^*, S^*)$ for each node v_i in this path. We will argue that $r_1 \geq 0$ and $r_{i+1} \geq r_i + 2$. It will immediately follow that $r_d \geq 2d-2$.

By Definition 4, $r_0 = r(v_0, G_N^*, S^*) \geq 0$. Now let us consider the edge $(v_i, v_{i+1}) \in E(G - {}_E S)$. We know that $(v_i, v_{i+1}) \notin S_{\text{Dest}} \cap S_{\text{ExtDisr}} \cap S_{\text{IntDisr}}$ and there must be a corresponding edge $(u_i^*, u_{i+1}^*) \in E(G_N^*)$. Such an edge (v_i, v_{i+1}) would have been deleted if u_{i+1}^* were destroyed, u_i^* were internally disrespectful, u_{i+1}^* were externally disrespectful, or either v_{i+1} or v_i were spoiled.

Therefore, the node $u_{i+1}^* \in \text{IN}(\text{SC}_{v_{i+1}})$ must have the following properties: First, $u_{i+1}^* \notin S^*$, i.e. it has not been removed/destroyed. Second, u_{i+1}^*'s parent, u_i^* is not in the set $\text{IntDisr}(G_N^*, S^*)$ i.e., u_i^* is not internally disrespectful. Third, $u_{i+1}^* \notin \text{ExtDisr}(G_N^*, S^*)$ i.e., u_{i+1}^* is not externally disrespectful.

Recall that the four properties of an externally disrespectful node u_{i+1}^* in a superconcentrator $\text{SC}v_{i+1}$, with parent u_i^* in a superconcentrator SC_{v_i} are: (1) u_{i+1}^* is not destroyed (2) $r(v_{i+1}, G_N^*, S^*) \leq r(v_i, G_N^*, S^*) + 1$ (3) u_{i+1}^*'s parent,

u_i^*, is not internally disrespectful (4) neither v_{i+1} nor v_i are spoiled. We know that u_{i+1}^* satisfies both the first and third properties, so if u_{i+1}^* satisfies the second property — that $r(v_{i+1}, G_N^*, S^*) \leq r(v_i, G_N^*, S^*) + 1$ — then $u_{i+1}^* \in$ would be externally disrespectful. Yet we know this is not the case, since u_{i+1}^* is not externally disrespectful. Thus, we must have $r_{i+1} = r(v_{i+1}, G_N^*, S^*) \geq r(v_i, G_N^*, S^*) + 2 = 2 + r_i$

Thus, by induction, we know that $r(v_d, G_N^*, S^*) \geq 2d - 2$. By Definition 4, we know that if $r(v_d, G_N^*, S^*) \geq 2d-2$, there exists a node $u'_d \in \text{IN}(\text{SC}(v_d))$ with depth $\texttt{depth}(v_d^*, G_N^* - _V S^*) \geq r(v_d, G_N^*, S^*) \geq 2d - 2$. Thus, there must be a path of length $2d - 2$ ending at u'_d in G_N^*. □

Now, we connect this to depth-robustness.

Theorem 5. *Given a graph G_N which is (e, d)-*edge *depth robust, a family of superconcentrators* SC*, and families $I = \{\texttt{in}_v\}_{v \in V}$ and $O = \{\texttt{out}_v\}_{v \in V}$ of one-to-one functions, let $G_N^* = \text{TRANSFORM}(G_N, \text{SC}, I, O)$ be the transformed graph constructed as in Definition 3. G_N^* is $(e/3, 2d-1)$-*node *depth robust for all $d > 1$.*

Proof. Suppose, for the sake of contradiction, that G_N^* is not $(e/3, 2d-1)$-*node* depth robust. Let S^* be the set of at most $e/3$ nodes removed from G_N^* such that $\texttt{depth}(G_N^* - _V S^*) < 2d$. Let $S = \text{NTOE}(G_N^*, S^*)$ be a set of edges in G_N as in Definition 10. We know from Lemma 4 that $|S| \leq 3|S^*|$.

Consider the graph $G_N - _E S$. Since $|S| \leq e$, and G_N is (e, d)-edge depth robust, $\texttt{depth}(G - _E S) \geq d$. Consider some such node v_d with $\texttt{depth}(v_d, G - _E S) = d$. There must be a path $P(v_d) = v_1, v_2, \ldots, v_d$ of length d ending in that node v_d. Thus, by Lemma 5, $G_N^* - _V S^*$ contains a path of length $2d - 2$ ending at a node $u'_d \in \text{IN}(\text{SC}_{v_d})$ with $r(v_d, G_N^*, S^*) \geq 2d - 2$. Because we know $d > 1$, there must be some edge (u'_{d-1}, u'_d) leading into the final node of P', which means SC_{v_d} is not spoiled (as all edges leading into a spoiled superconcentrator are removed). This, in turn, means that $\gamma_{v_d}(S^*) < \frac{2}{3}$ i.e., fewer than $\frac{2}{3}|\text{IN}(\text{SC}_{v_d})|$ nodes were removed from the superconcentrator SC_{v_d}. We know by Lemma 2 that there are at most $|\text{INTDISR}(v_d, G_N^*, S^*)| \leq |S_v^*| = \gamma_{v_d}(S^*)|\text{OUT}(\text{SC}_{v_d})| < \frac{2}{3}|\text{OUT}(\text{SC}_{v_d})|$ internally disrespectful nodes in $\text{OUT}(\text{SC}_{v_d})$. Define the set $\text{GOOD}_{\text{OUT}}(\text{SC}_{v_d}) = \text{OUT}(\text{SC}_{v_d}) \setminus \text{INTDISR}(v_d, G_N^*, S^*)$ and note that $|\text{GOOD}_{\text{OUT}}(\text{SC}_{v_d})| \geq |\text{OUT}(\text{SC}_{v_d})| - \frac{2}{3}|\text{IN}(\text{SC}_{v_d})| = \frac{1}{3}|\text{IN}(\text{SC}_{v_d})|$. In particular, the set $\text{GOOD}_{\text{OUT}}(\text{SC}_{v_d})$ is non-empty. Now observe that, by definition, for any $x \in \text{GOOD}_{\text{OUT}}(\text{SC}_{v_d})$ we have $\texttt{depth}(v, G_N^* - _V S^*) \geq 1 + r(v_d, G_N^*, S^*) \geq 2d - 1$ — see Definition 8. Thus, $\texttt{depth}(G_N^* - _V S^*) \geq 2d$ which contradicts our original assumption. It follows that G_N^* is $(e/3, 2d-1)$-node depth robust. □

B Fractional Depth-Robustness

Beyond just simple depth-robustness, we also consider the more nuanced idea of *fractional* depth-robustness. Intuitively, a graph G is fractionally node-depth-robust if, when a set of nodes (of size at most e) is removed, the remaining graph has a *large fraction* f_2 (respectively f_1) of its nodes end (resp. start) paths of depth at least d — we count the number of sinks, ending a long path

(resp sources, starting a long path), and divide by the total number of nodes, $|V(G)|$. We will formally define this in Definition 11. Generally, we want to have $f_2, f_1 \in \Omega(1)$, so that a constant fraction of nodes are sources and sinks of long paths.

In Lemma 5 (Section A) we showed that for any $S^* \subseteq V(G_N^*)$ of nodes in G_N^* that the set $S = \text{NTOE}(S) \subseteq E(G_N)$ has size $|S| \leq 3N$ and that any path P_v in $G_N^- \text{NTOE}(S)$ ending at node v corresponds to a similar length P^* in G_N^* ending at some node in $\text{OUT}(\text{SC}(v))$. Thus, it seems like we should immediately be able to conclude that if G_N is $(e, d, 0, f_2)$ fractionally edge depth robust that G_N^* is $(\Omega(e), \Omega(d), 0, f_2)$ fractionally node depth robust. Unfortunately, this conclusion does not quite hold because f_2 represents the fraction of nodes in the graph that have depth at least d and the total number of nodes in G_N^* may have increased significantly e.g., if $|E(G_N))| = \Omega(N \log N)$ then we can only conclude that G_N^* is $(e' = \Omega(e), d' = \Omega(d), 0, f_2')$ fractionally node depth robust with $f_2' = f_2 \cdot \Omega\left(\frac{N}{|E(G_N)|}\right)$. While this guarantee may still be slightly meaningful, it is not tight and is less useful for applications to dMHFs [11] and Proofs of Space [22,8] i.e., an attacker who starts with just e' pebbles on the graph may still be able to respond to a random challenge with probability $1 - o(1)$. Instead, in this section we will prove a stronger bound on fractional node-depth-robustness, Theorem 6.

Theorem 6. *Given a graph G_N which is (e, d, f_1, f_2) fractionally edge depth robust, a family of superconcentrators $\text{SC} = \{\text{SC}(i)\}_{i=1}^{\texttt{mdeg}(G_N)}$, and families of one-to-one functions $I = \{\texttt{in}_v : \texttt{parents}(v, G_n) \to \text{IN}(\text{SC}_v)\}_{v \in G_n}$ and $O = \{\texttt{out}_v : \texttt{children}(v, G_n) \to \text{OUT}(\text{SC}_v)\}_{v \in G_n}$. Let $G_N^* = \text{TRANSFORM}(G_N, \text{SC}, I, O)$ be the corresponding graph constructed as in Definition 3 and let $\beta = \frac{1}{3} \cdot \frac{|V(G_N)|}{|V(G_N^*)|} \cdot \min_{v \in V(G_N)} \texttt{mdeg}(v)$. G_N^* is $(e/3, 2d-1, f_1 \cdot \beta, f_2 \cdot \beta)$ fractionally node depth robust for all $d > 1$.*

Applying Theorem 6 to G_N we obtain a graph G_N^* which has $N' = O(N \log N)$ nodes and constant indegree and is (e', d', f_1', f_2') fractionally node depth robust for $e' = \Omega(N')$, $d' = \Omega(N'^{1-\epsilon}/\log N)$, $f_1, f_2 = \Omega(N)$.

B.1 Reversibility of Transformation

We first show that our transformation is reversible i.e., if you reverse the transformed graph you still obtain a node depth robust graph. This will allow us to focus on proving the our transformed graph is $(e, d, 0, f_2)$ node depth-robust, and then we will be able to invoke reversibility to conclude that the transformed graph is also $(e, d, f_2, 0)$ node depth-robust.

Fact 7 *If SC is a superconcentrator with n inputs I and n outputs O then $\texttt{Rev}(\text{SC})$ is also a superconcentrator with n inputs O and n outputs I.*

Proof. Consider any subsets $A \subset O$ and $B \subseteq I$ with $|B| = |A|$. We want to show that there are $|B|$ vertex disjoint paths from A to B in $\texttt{Rev}(\text{SC})$. Because we know SC is a superconcentrator we can find $|B|$ vertex disjoint paths $P_1, \ldots, P_{|B|}$ in SC

where each path P_i starts in B and ends in A. For each path P_i we let R_i denote the corresponding reversed path in $\texttt{Rev}(\text{SC})$. Observe that the paths $R_1, \ldots, R_{|B|}$ are still vertex disjoint and that each path now starts in A and ends in B. Thus, $\texttt{Rev}(\text{SC})$ is a superconcentrator. □

Lemma 6. *Given a graph $G_N = (V, E)$, a family of superconcentrators $\text{SC} = \{\text{SC}(i)\}_{i=1}^{\texttt{mdeg}(G_N)}$ we let SC' be the family of superconcentrators given by reversing SC, i.e., $\text{SC}' \doteq \{\text{SC}'(i) = \textit{Rev}(\text{SC}(i))\}_{i=1}^{\texttt{mdeg}(G_N)}$. For any families of one-to-one functions $I = \{\texttt{in}_v : \texttt{parents}(v, G_n) \to \text{IN}(\text{SC}_v)\}_{v \in G_n}$ and $O = \{\texttt{out}_v : \texttt{children}(v, G_n) \to \text{OUT}(\text{SC}_v)\}_{v \in G_n}$, let I' and O' be the families of one-to-one functions given by $I' \doteq \{\texttt{in}'_v = \texttt{out}_v\}_{v \in G_n}$ and $O' \doteq \{\texttt{out}'_v = \texttt{in}_v\}_{v \in G_n}$. Then $\textit{Rev}(\text{TRANSFORM}(G_N, \text{SC}, I, O)) = \text{TRANSFORM}(\textit{Rev}(G_N), \text{SC}', I', O')$.*

Proof. The transformation $\text{TRANSFORM}(G_N, \text{SC}, I, O)$ works by replacing each node $v \in V_N$ with a superconcentrator SC_v which is a fresh copy of $\text{SC}(\texttt{mdeg}(v))$ and then connecting adding edges between the superconcentrators. Because $\texttt{Rev}(G_N)$ and G_N have the same set of nodes $\text{TRANSFORM}(\texttt{Rev}(G_N), \text{SC}', I', O')$ will simply replace each node with a fresh copy of $\text{SC}'(\texttt{mdeg}(v))) = \texttt{Rev}(\text{SC}(\texttt{mdeg}(v)))$ and then add edges between the superconcentrators. Let $H_N = (V_N, E'_N)$ where $W_N = \bigcup_{v \in V_N} V(\text{SC}_v)$ denotes the graph $\texttt{Rev}(\text{TRANSFORM}(G_N, \text{SC}, I, O))$. Similarly, let $H'_N = (W_N, E''_N)$ denote the graph $\text{TRANSFORM}(\texttt{Rev}(G_N), \text{SC}', I', O')$ — note that the set of nodes for both graphs H_N and H'_N are identical.

We first consider edges internal to the superconcentrators and observe that replacing each node v with a superconcentrator SC_v then reversing these edges is equivalent to replacing node v with a fresh copy of $\text{SC}'(\texttt{mdeg}(v)) = \texttt{Rev}(\text{SC}(\texttt{mdeg}(v)))$. Thus, for any $v \in V_N$ and any $x, y \in V(\text{SC}_v)$ we have $(x, y) \in E'_N$ if and only if $(x, y) \in E''_N$.

We also need to consider the edges between the superconcentrators as determined by the functions $\texttt{in}/\texttt{in}'$ and $\texttt{out}/\texttt{out}'$. Consider any edge $(u, v) \in E_N$ in the original graph G_N. To construct $\texttt{Rev}(\text{TRANSFORM}(G_N, \text{SC}, I, O))$, we first add an edge from $\texttt{out}_u(v)$ (which is in $\text{OUT}(\text{SC}_u)$) to $\texttt{in}_v(u)$ (which is in $\text{IN}(\text{SC}_v)$). After reversal, this means we have an edge from $\texttt{in}_v(u)$ to $\texttt{out}_u(v)$. On the other hand, to construct $\text{TRANSFORM}(\texttt{Rev}(G_N), \text{SC}', I', O')$, we first reverse the graph G_N, resulting in an edge $(v, u) \in \texttt{Rev}(E_N)$. We then add a corresponding edge from $\texttt{out}'_v(u) = \texttt{in}_v(u)$ to $\texttt{in}'_u(v) = \texttt{out}_u(v)$. Thus, for all $u \neq v \in V_N$ and all $x \in \text{SC}_u$ and all $y \in SC_v$ we have $(x, y) \in E'_N$ if and only if $(x, y) \in E''_N$. Therefore, $E'_N = E''_N$. It follows that the graphs $\texttt{Rev}(\text{TRANSFORM}(G_N, \text{SC}, I, O))$ and $\text{TRANSFORM}(\texttt{Rev}(G_N), \text{SC}', I', O')$ are indeed identical. □

B.2 Defining Fractional Depth-Robustness

We can extend our ideas of edge and node depth-robustness to those of *fractional* depth-robustness. Intuitively, a graph is fractionally depth robust if, when any set S of at most $|S| \leq e$ of vertices is removed, the remaining graph not only has a node of depth d, but *many* (at least $f_2 N$) nodes in $G - S$ are the *endpoint* of a direct path of length d, and *many* (at least $f_1 N$) nodes in $G - S$

are the source of some directed path of length d. We note that this does not necessarily imply that there are many *node disjoint* paths of length d.

Definition 11. *Given a DAG G we say that G is (e, d, f_1, f_2)-depth robust if, for all $S \subseteq V(G)$ with $|S| \leq e$, we have $|\{v \in V(G) : \texttt{depth}(v, G-_V S) \geq d\}| \geq f_2 \cdot |V(G)|$ and $|\{v \in V(\texttt{Rev}(G)) : \texttt{depth}(v, \texttt{Rev}(G)-_V S) \geq d\}| \geq f_1 \cdot |V(G)|$.*

This definition generalizes the notion of fractional depth-robustness [14] to include fractional source depth-robustness as well as fractional sink depth-robustness. In particular, the notion of (e, d, f)-fractional depth robustness from [14] is equivalent to $(e, d, 0, f)$-depth robustness in Definition 11 i.e., setting $f_1 = 0$ and $f_2 = f$. We can also define fractional edge-depth-robustness, where a set of edges, rather than a set of nodes, is removed.

Definition 12. *Given a DAG $G = (V, E)$ we say that G is (e, d, f_1, f_2)-edge depth robust if, for all $S \subseteq E$ with $|S| \leq e$, we have $|\{v \in V : \texttt{depth}(v, G-_E S) \geq d\}| \geq f_2 \cdot |V(G)|$ and for all $S \subseteq \texttt{Rev}(E)$ with $|S| \leq e$ we have $|\{v \in V(\texttt{Rev}(G)) : \texttt{depth}(v, \texttt{Rev}(G-_E S)) \geq d\}| \geq f_1 \cdot |V(G)|$.*

We also note that a graph G is (e, d, f_1, f_2)-depth robust if and only if the reverse graph $\texttt{Rev}(G)$ is (e, d, f_2, f_1)-depth robust. Similarly, G is (e, d, f_1, f_2)-edge depth robust if and only if the reverse graph $\texttt{Rev}(G)$ is (e, d, f_2, f_1)-edge depth robust.

B.3 Our Transform Preserves Fractional Depth-Robustness

Just as we showed that our construction converts any edge-depth robust graph G into a node-depth robust graph G^*_N, we show that it converts any fractional edge-depth robust graph G into a fractional node-depth robust graph G^*_N. Intuitively, if a set of nodes is removed from G^*_N, Definition 10 gives a corresponding set of edges to remove from G. We know from Definition 12 that there are many sink nodes in G after removing a set of edges. Each of these sink nodes has high depth, which corresponds to a long path of superconcentrators with increasing r values. For each of these paths, the final superconcentrator's high r value implies it has a node with large depth.

Lemma 7. *Given a graph G_N which is $(e, d, 0, f_2)$ fractionally edge depth robust, let $G^*_N = \textsc{Transform}(G_N, \text{SC}, I, O)$ be the transformed graph constructed as in Definition 3 and let $\beta = \frac{1}{3} \cdot \frac{|V(G_N)|}{|V(G^*_N)|} \cdot \min_{v \in V(G_N)} |\textsc{out}(\text{SC}_{v_{i,d}})|$. G^*_N is $(e/3, 2d-1, 0, f_2 \cdot \beta)$ fractionally node depth robust for all $d > 1$.*

Proof. Consider any set $S^* \subseteq V(G^*_N)$ of at most $e/3$ nodes to be removed from G^*_N. Let $S = \textsc{NtoE}(G^*_N, S^*)$ be the set of edges in G_N given by applying Definition 10 to S^*; we know by Lemma 4 that $|S| \leq 3|S^*|$.

Consider the graph $G_N-_E S$. We know that $|S| \leq e$, and since G_N is $(e, d, 0, f_2)$ fractionally edge depth robust, there exists a set $D = \{v \in V(G_N) : \texttt{depth}(v, G_N-_E S) \geq d\}$ such that $|D| \geq f_2 \cdot V(G_N)$. For each node $v_{i,d} \in D$, there must be a path $P(v_{i,d}) = v_{i,1}, v_{i,2}, \ldots, v_{i,d}$ in $G_N-_E S$. By Lemma 5, for each $v_{i,d}$, there

is a path of length at $2d-2$ ending at some node $u'_{i,d} \in \text{IN}(\text{SC}_{v_{i,d}})$. Let D' be the set $D' = \bigcup_{v_{i,d} \in D} \{u'_{i,d}\}$ of each these long path nodes in G^*_N. Since each node $v_{i,d}$ must be unique, each of their superconcentrators, and thus all nodes in the input set of their superconcentrators, must also be unique. Thus, $|D'| \geq |D| \geq f_2 \cdot |V(G_N)|$.

Now, consider each of the superconcentrators containing these nodes. Since each $u'_{i,d}$ has depth at least $2d-2$, and we know $d > 1$, we know there must be an edge $\left(u'_{i-1,d}, u'_{i,d}\right)$ leading into it, and that $\text{SC}_{v_{i,d}}$ must therefore not be spoiled (as all edges leading into a spoiled superconcentrator are removed). This, in turn, means that $\gamma_{v_{i,d}}(S^*) < \frac{2}{3}$. We know by our application of Lemma 5 that $r(v_{i,d}, G^*_N, S^*) \geq 2d-2$ Now, consider the set of output nodes from $\text{SC}_{v_{i,d}}$, $\text{OUT}(\text{SC}_{v_{i,d}})$. By Lemma 2 that there are at most $|\text{INTDISR}(v_{i,d}, G^*_N, S^*)| \leq |S^*_v| = \gamma_{v_{i,d}}(S^*)|\text{OUT}(\text{SC}_{v_{i,d}})| < \frac{2}{3}|\text{OUT}(\text{SC}_{v_{i,d}})|$ externally disrespectful nodes in $\text{OUT}(\text{SC}_{v_{i,d}})$. Define the set $\text{GOOD}_{\text{OUT}}(\text{SC}_{v_{i,d}}) = \text{OUT}(\text{SC}_{v_{i,d}}) \backslash \text{INTDISR}(v_{i,d}, G^*_N, S^*)$ and note that $\left|\text{GOOD}_{\text{OUT}}(\text{SC}_{v_{i,d}})\right| \geq \left|\text{OUT}(\text{SC}_{v_{i,d}})\right| - \frac{2}{3}\left|\text{IN}(\text{SC}_{v_{i,d}})\right| = \frac{1}{3}\left|\text{IN}(\text{SC}_{v_{i,d}})\right|$. Now observe that, by definition, for any $x \in \text{GOOD}_{\text{OUT}}(\text{SC}_{v_{i,d}})$ we have $\texttt{depth}(v, G^*_N - _V S^*) \geq 1 + r(v_d, G^*_N, S^*) \geq 2d-1$ — see Definition 8. Thus, there remains a set $\text{GOOD}_{\text{OUT}}(\text{SC}_{v_{i,d}})$ of at least $\frac{1}{3}|\text{OUT}(\text{SC}_{v_{i,d}})|$ internally respectful nodes in $\text{OUT}(\text{SC}_{v_{i,d}})$ with $\forall u \in \text{GOOD}_{\text{OUT}}(\text{SC}_{v_{i,d}}) \texttt{depth}(u, G^*_N - _V S^*) \geq r(v_{i,d}, G^*_N, S^*) + 1 \geq 2d-1$.

Now further define the set of all good outputs $\text{GOOD} = \bigcup_{v_{i,d} \in D} \text{GOOD}_{\text{OUT}}(\text{SC}_{v_{i,d}})$. Since each superconcentrator's output set is distinct, $|\text{GOOD}| = \Sigma_{v_{i,d} \in D}|\text{GOOD}_{\text{OUT}}(\text{SC}_{v_{i,d}})| \geq \Sigma_{v_{i,d} \in D} \frac{1}{3}|\text{OUT}(\text{SC}_{v_{i,d}})| \geq f_2 \cdot |V(G_N)| \cdot \min_{v_{i,d} \in V(G_N)} \frac{1}{3}|\text{OUT}(\text{SC}_{v_{i,d}})|$. This gives us a set GOOD of at least $\frac{1}{3} f_2 \cdot |V(G_N)| \cdot \min_{v_{i,d} \in V(G_N)} |\text{OUT}(\text{SC}_{v_{i,d}})|$ distinct nodes in G^*_N with depth at least $2d-1$. Thus, G^*_N is $(e/3, 2d-1, 0, f_2 \cdot \beta)$ fractionally node depth robust. □

Reminder of Theorem 6. *Given a graph G_N which is (e, d, f_1, f_2) fractionally edge depth robust, a family of superconcentrators $\text{SC} = \{\text{SC}(i)\}_{i=1}^{\texttt{mdeg}(G_N)}$, and families of one-to-one functions $I = \{\texttt{in}_v : \texttt{parents}(v, G_n) \to \text{IN}(\text{SC}_v)\}_{v \in G_n}$ and $O = \{\texttt{out}_v : \texttt{children}(v, G_n) \to \text{OUT}(\text{SC}_v)\}_{v \in G_n}$. Let $G^*_N = \text{TRANSFORM}(G_N, \text{SC}, I, O)$ be the corresponding graph constructed as in Definition 3 and let $\beta = \frac{1}{3} \cdot \frac{|V(G_N)|}{|V(G^*_N)|} \cdot \min_{v \in V(G_N)} \texttt{mdeg}(v)$. G^*_N is $(e/3, 2d-1, f_1 \cdot \beta, f_2 \cdot \beta)$ fractionally node depth robust for all $d > 1$.*

Proof of Theorem 6. Since G_N is (e, d, f_1, f_2) fractionally edge depth robust, it is also $(e, d, 0 f_2)$ fractionally edge depth robust. We know by Lemma 7 that G^*_N is therefore $(e/3, 2d-1, 0, f_2 \cdot \beta)$ fractionally node depth robust for all $d > 1$.

Consider also the reverse graph, $\texttt{Rev}(G_N)$. We know $\texttt{Rev}(G_N)$ is $(e, d, 0, f_1)$ fractionally edge depth robust, and by applying Lemma 7 that $\text{TRANSFORM}(\texttt{Rev}(G_N), \text{SC}', I', O')$ is $(e/3, 2d-1, 0, f_1 \cdot \beta)$ fractionally node depth robust for any family of superconcentrators SC' and families of one-to-one functions I' and O' (such that we still have $\beta = \frac{|V(G_N)|}{|V(\text{TRANSFORM}(\texttt{Rev}(G_N), \text{SC}', I', O'))|}$). Further, if we let $\text{SC}' = \{\text{SC}'(i) = \texttt{Rev}(\text{SC}(i))\}_{i=1}^{\texttt{mdeg}(v)}$, $I' = \{\texttt{in}'_v = \texttt{out}_v\}_{v \in G_n}$ and $O' = \{\texttt{out}'_v = \texttt{in}_v\}_{v \in G_n}$, we know by Lemma 6 that $\text{TRANSFORM}(\texttt{Rev}(G_N), \text{SC}', I', O') = \texttt{Rev}(\text{TRANSFORM}(G_N, \text{SC}, I, O))$. Thus, the reverse of G^*_N, $\texttt{Rev}(\text{TRANSFORM}(G_N, \text{SC}, I, O))$,

is $(e/3, 2d-1, 0, f_1 \cdot \beta)$ fractionally node depth robust for all $d > 1$. This means $\texttt{Rev}(\texttt{Rev}(\textsc{Transform}(G_N, \textsc{sc}, I, O))) = G^*_N$ is $(e/3, 2d-1, f_1 \cdot \beta, 0)$ fractionally node depth robust for all $d > 1$.

Since G^*_N is both $(e/3, 2d, 0, f_2 \cdot \beta)$ and $(e/3, 2d, f_1 \cdot \beta, 0)$ fractionally node depth robust for all $d > 1$, it follows from Definition 11 that G^*_N is $(e/3, 2d-1, f_1 \cdot \beta, f_2 \cdot \beta)$ fractionally node depth robust for all $d > 1$. □

As a concrete example, the Grates graph $G^{\texttt{Grates}}_n$ has $O(N \log N)$ edges and is (e, d, f_1, f_2)-fractionally edge depth for $e = \Omega(N)$, $d = \Omega(N^{1-\epsilon})$ and $f_1, f_2 = \Omega(1)$ [25]. For this particular graph we have $\min_v \min_v \texttt{mdeg}(v) = \Omega(\log N)$ so when we apply the transformation from Theorem 6 we have $\beta = \frac{N}{3N \log N} \min_v \texttt{mdeg}(v) = \Omega(1)$. In particular, the graph G^*_N has $N' = O(N \log N)$ nodes and constant indegree and is (e', d', f'_1, f'_2) fractionally node depth robust for $e' = \Omega(N')$, $d' = \Omega(N'^{1-\epsilon}/\log N)$, $f_1, f_2 = \Omega(N)$. Thus, if we start with e' pebbles on the graph G^*_N and are challenged to place a pebble on a randomly selected node it will take at least d' rounds to do this with probability $f'_2 = \Omega(1)$. This property is useful for designing data-dependent Memory Hard Functions (dMHFs) [11] as well as Proofs of Space [22,8]. In particular, a data-dependent MHF has dynamic edges which are not known a priori because they are are data-dependent. Intuitively, any evaluation algorithm either must store e' labels or risk an unlucky challenge which will delay further progress for at least d' rounds. For a Proof of Space we can ensure that any cheating prover that does not store at least e' labels will (whp) not be able to respond to k random challenges without performing d' rounds of sequential computation. Thus, by demanding that a prover respond quickly to all k challenges a cheating prover will be caught with high probability.

C Proofs of Section 3.3

We have shown previously that our transformation converts any (e, d) edge-depth robust graph into an $(e/3, 2d-1)$ node-depth robust graph. However, as long as $N = |V(G_N)| = o(|E(G_N)|)$, as we would expect in any useful construction of an edge-depth robust graph, we will necessarily have $d = o(|V(G^*_N)|)$ because the transformation increases the number of vertices in G^*_N such that $|V(G^*_N)| \approx |E(G_N)|$. By contrast, there are direct node-depth robust constructions to achieve (e, d)-node depth robustness with $e = \Omega(N/\log N)$ and $d = \Omega(N)$ e.g., see [3,4]. At first glance it seems impossible to achieve a similar result using our transformation. In this section we show that if the superconcentrator family is "d'-perfectly connected" then our transformation can convert our (e, d) edge-depth robust graph into a $(e/3, (d-e) \cdot d')$ node depth robust graph. In particular, if $|E(G_N)| = O(N \log N)$, $d' = \Omega(\log N)$, and $d = \Omega(N)$ then the final depth parameter $d \cdot d' = \Omega(N \log N) = \Omega(|V(G^*_N)|)$ is linear in the size of the final graph graph G^*_N.

This results in robustness that is smaller proportional to the size of the graph. In this section, we will show that, for particular families of superconcentrators and (relatively) small values of e, our transformation can be shown to have superior node-depth robustness, reducing this proportional loss. To do so, we

must first define a class of d'-perfectly connected superconcentrators. Intuitively, a superconcentrator is d'-perfectly connected if there is a *long* path from all inputs to all outputs. Note that this only requires a single such long path — for sets of multiple inputs and outputs, the vertex disjoint paths guaranteed by Definition 1 may all be short.

Definition 13. *Given a DAG $G = (V, E)$ with n input nodes $\text{IN}(G) \subseteq V$ and n output nodes $\text{OUT}(G) \subseteq V$ we say that G and a depth parameter $d' \geq 1$ we say that G is d'-perfectly connected if $\forall u \in \text{IN}(G))$ and $\forall v \in \text{OUT}(G)$, there exists a path $P_{u,v} = (u, p_1, p_2, \ldots, v)$ in G from u to v of length at least d', i.e. $|P| \geq d'$.*

Note that d'-perfect connectivity does not imply that G is a superconcentrator. Similarly, a superconcentrator $\text{SC}(n)$ does not necessarily satisfy the property of d'-perfect connectivity for large depth parameters $d' = \Omega(n)$. We will be interested in constructing superconcentrators which also satisfy the property of d'-perfect connectivity i.e., a d'-perfectly connected superconcentrator. We will use the notation $\text{SC}^{depth}(n, d')$ to refer to a superconcentrator which satisfies both properties. Given a family of superconcentrators $\text{SC} = \{\text{SC}(n)\}_{n=1}^{\infty}$ we use $\text{SC}^{depth} = \{\text{SC}^{depth}(n, d')\}$ to refer to the new transformed family of superconcentrators which all satisfy d'-perfect connectivity.

Given any n-superconcentrator $\text{SC}(n)$, we can construct a d'-perfectly connected superconcentrator as depicted in Figure 4. New nodes and edges are depicted in red. We formalize this construction in Definition 14.

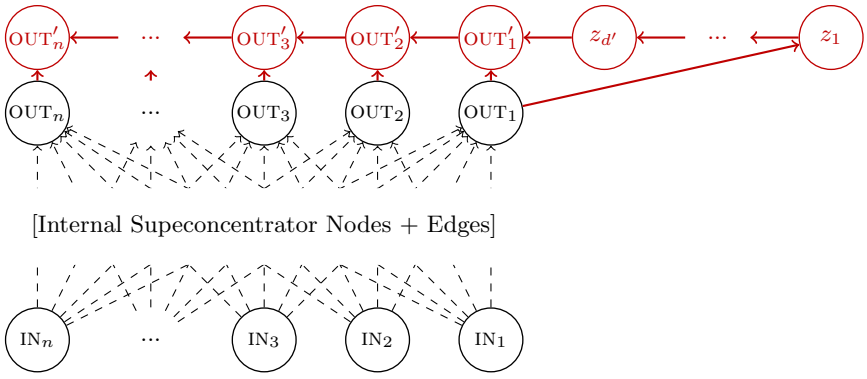


Fig. 4: $\text{SC}^{depth}(n, d')$, a d'-perfectly connected superconcentrator

Definition 14. *Given a n-superconcentrator $\text{SC}(n)$, and a depth parameter d'. Let $\text{SC}^{depth}(n, d')$ be constructed as follows: Define $\text{NEW} = \{z_1, \ldots, z_{d'}\} \cup \{\text{OUT}'_1, \ldots, \text{OUT}'_n\}$ by adding d' new internal nodes $z_1 \ldots z_{d'}$ and n new output nodes $\text{OUT}'_1, \ldots, \text{OUT}'_n$. Adding edges $(\text{OUT}_i, \text{OUT}'_i)$ for all $i \leq n$ will ensure that the new graph $\text{SC}^{depth}(n, d')$ is still a n-superconcentrator. To ensure that the graph is d'-perfectly connected we will also include edges (1) (z_i, z_{i+1}) for $i < d'$, (2) $(\text{OUT}'_i, \text{OUT}'_{i+1})$ for $i < n$,*

(3) (OUT_1, z_1)*, and (4)* $(z_{d'}, \text{OUT}'_1)$*. Formally,* $V\left(\text{SC}^{depth}(n,d')\right) = V(\text{SC}(n)) \cup$ NEW *and* $E\left(\text{SC}^{depth}(n,d')\right) = E\left(\text{SC}(n)\right) \cup \{\left(\text{OUT}'_i, \text{OUT}'_{i+1}\right)\}_{i=1}^{n-1} \cup \{(z_i, z_{i+1})\}_{i=1}^{n} \cup \{(\text{OUT}_1, z_1)\} \cup \{(z'_d, \text{OUT}'_1)\} \cup \{(\text{OUT}_i, \text{OUT}'_i)\}_{i=1}^{n}$*. Let us define the new input and output sets* $\text{OUT}(\text{SC}^{depth}(n,d')) = \{\text{OUT}'_1, \ldots, \text{OUT}'_n\}$ *and* $\text{IN}(\text{SC}^{depth}(n,d')) = \text{IN}(\text{SC}(n))$.

The graph $\text{SC}^{depth}(n,d')$ adds $n+d'$ new nodes and $2n+d'-1$ new edges to the original superconcentrator $\text{SC}(n)$.

Claim 8 *Given any superconcentrator* $\text{SC}(n)$ *and depth parameter* d'*, let* $\text{SC}^{depth}(n,d')$ *be constructed as in Definition 14. Then the graph* $\text{SC}^{depth}(n,d')$ *is an* n*-superconcentrator, is* d'*-perfectly connected, has* $\texttt{indeg}(\text{SC}^{depth}(n,d')) = \max(\texttt{indeg}(\text{SC}(n), 2)$*, and has* $|V(\text{SC}^{depth}(n,d'))| = |V(\text{SC}(n))| + n + d'$.

Proof. We first verify that the graph $\text{SC}^{depth}(n,d')$ is still a superconcentrator. Fix any $S_{in} \subseteq \text{IN}$ and any $S'_{out} \subseteq \text{OUT}'$ of equal size $|S_{in}| = |S_{out}| = k$. We will let $S_{out} = \{\text{OUT}_i : \text{OUT}'_i \in S'_{out}\}$ denote the corresponding set of outputs from the original superconcentrator $\text{SC}(n)$. By definition the original superconcentrator $\text{SC}(n)$ contains k node disjoint paths $P_1, \ldots, P_k$ connecting S_{in} to S_{out}. Using the edges $(\text{OUT}_i, \text{OUT}'_i)$ we can extend each path to obtain k nodes disjoint paths $P'_1, \ldots, P'_k$ connecting S_{in} to S_{out}.
To see that the graph is d'-perfectly connected fix an arbitrary input/output pair IN_i and OUT'_j. Since $\text{SC}(n)$ is a superconcentrator there is a path from IN_i to OUT_1, and by construction, there is a directed path $\text{OUT}_1, z_1, \ldots, z_{d'}, \ldots, \text{OUT}_{j'}$ from OUT_1 to OUT'_j of length $\geq d'$. □ It follows from Definition 14 that we only add nodes of indegree 2, and that we add $d' + N$ new nodes. □

Such d'-perfectly connected superconcentrators can yield an advantage when used in our transformation — so long as a superconcentrator does not have any nodes removed by a reducing set, we know that it will increase the depth of paths through it not just by 1, but by d'. This means that if e is suitably small, and a large number of superconcentrators must be untouched, we can make stronger guarantees about the depth of the remaining paths.

We still rely on our original definitions of the SPOILED and DEST nodes, Definitions 5 and 7. Since our original set INTDISR was based off an underlying assumption that each superconcentrator must have at least one edge between input and output, we further refine it to the set of d-perfectly connected internally disrespectful nodes, INTDISRHD, formalized in Definition 15.

Definition 15. *Given a graph* G_N*, a family of superconcentrators* SC^{depth} *with perfect connectivity at depth* d'*, and families* $I = \{\texttt{in}_v\}_{v \in V}$ *and* $O = \{\texttt{out}_v\}_{v \in V}$ *of one-to-one functions, let* $G^*_N = \text{TRANSFORM}(G_N, \text{SC}^{depth}, I, O)$ *be the transformed graph constructed as in Definition 3. Given a superconcentrator* SC_v^{depth} *in* G^*_N *and a set* S^* *of nodes removed from* G^*_N*, we define the set* $\text{INTDISRHD}(v, G^*_N, S^*)$ *to be: (1) the set of all output nodes which do not have depth at least* $r(v, G^*_N, S^*) + 1$ *if any node was removed from* SC_v*, i.e.* $\text{INTDISRHD}(v, G^*_N, S^*) \doteq \{x \in \text{OUT}(\text{SC}_v) \ : \ \texttt{depth}(x, G^*_N - _V S^*) < r(v, G^*_N, S^*) + 1\}$ *if* $|S^*_v| > 0$ *or (2) the set of all output nodes which do not*

have depth at least $r(v, G^*_N, S^*) + d'$ *if no node was removed from* SC_v*, i.e.* $\textsc{IntDisrHD}(v, G^*_N, S^*) \doteq \{x \in \textsc{out}(\text{SC}_v) : \texttt{depth}(x, G^*_N -_V S^*) < r(v, G^*_N, S^*) + d'\}$ *if* $|S^*_v| = 0$*. We further define the set* $\textsc{IntDisrHD}(G^*_N, S^*) = \bigcup_{v \in V(G_N)} \textsc{IntDisrHD}(v, G^*_N, S^*)$ *of all internally disrespectful nodes in the graph* G^*_N*.*

We also upper bound the number of internally disrespectful nodes in a given high-depth superconcentrator.

Lemma 8. *Given a graph* G_N*, a family of superconcentrators* SC^{depth} *with perfect connectivity at depth* d'*, and families* $I = \{\texttt{in}_v\}_{v \in V}$ *and* $O = \{\texttt{out}_v\}_{v \in V}$ *of one-to-one functions, let* $G^*_N = \textsc{Transform}(G_N, \text{SC}^{depth}, I, O)$ *be the transformed graph constructed as in Definition 3. Given a superconcentrator* SC^{depth}_v *in* G^*_N *and a set* S^* *of nodes removed from* G^*_N*,* $|\textsc{IntDisrHD}(v, G^*_N, S^*)| \leq |S^*_v|$

Proof. Suppose, for the sake of contradiction, that there are more than $|S^*_v|$ nodes in $\textsc{IntDisrHD}(v, G^*_N, S^*)$. Then there exists a set $B \subseteq \textsc{IntDisrHD}(v, G^*_N, S^*)$ of at least $|B| \geq |S^*_v| + 1$ internally disrespectful output vertices. Either $|S^*_V| > 0$ or $|S^*_v| = 0$. If $|S^*_V| > 0$, then Definition 15 is identical to Definition 8 and, by Lemma 2, we know $|\textsc{IntDisrHD}(v, G^*_N, S^*)| \leq |S^*_v|$. On the other hand, if $|S^*_v| = 0$, then no nodes were removed from SC^{depth}_v. Since it has perfect connectivity at depth d', by Definition 15 there is a path from every input to every output of length at least d'. By the definition of r, Definition 4, there is at least one node a with depth $\texttt{depth}(a) \geq r(v, G^*_N, S^*)$, Thus, for every output there is a path (through a) of depth at least $r(v, G^*_N, S^*) + d'$. This means that no output node can be high-depth internally disrespectful. This contradiction means that $|\textsc{IntDisrHD}(v, G^*_N, S^*)| \leq |S^*_v|$. □

Similarly, we must refine our notion of externally-disrespectful nodes slightly in a d'-perfectly connected context. we define a new set, ExtDisrHD, of externally disrespectful nodes for d'-perfectly connected superconcentrators. Once again, a node is externally disrespectful if the r value of its superconcentrator does not sufficiently increase, but now that increase may need to be d' instead of 1.

Definition 16. *Given a graph* G_N*, a family of superconcentrators* SC^{depth} *with perfect connectivity at depth* d'*, and families* $I = \{\texttt{in}_v\}_{v \in V}$ *and* $O = \{\texttt{out}_v\}_{v \in V}$ *of one-to-one functions, let* $G^*_N = \textsc{Transform}(G_N, \text{SC}^{depth}, I, O)$ *be the transformed graph constructed as in Definition 3. Given a superconcentrator* SC^{depth}_v *in* G^*_N *and a set* S^* *of nodes removed from* G^*_N*, for any node* $b \in \textsc{in}(\text{SC}^{depth}_v)$*, we say that* b *is "externally disrespectful" if* $\exists u \in V(G_N)$ *such that* $\texttt{In}_{\texttt{v}}(u) = b$ *(where* $\texttt{In}_v$ *is the function defined in Definition 3) and (1)* $b \notin \textsc{Dest}(v, G^*_N, S^*)$*, (2) Either* $r(v, G^*_N, S^*) \leq r(u, G^*_N, S^*) + 1$*, and* $|S^*_u| > 0$ *or* $r(v, G^*_N, S^*) \leq r(u, G^*_N, S^*) + d'$ *and* $|S^*_u| = 0$ *(3) The node* $a = \texttt{out}(u, v)$ *is not internally disrespectful, i.e.* $a \notin \textsc{IntDisrHD}(u, G^*_N, S^*)$*, and (4) neither* v *nor* u *is spoiled, i.e.* $\textsc{Spoiled}(v, G^*_N, S^*) = 0 \wedge \textsc{Spoiled}(u, G^*_N, S^*) = 0$*.*

We define the set $\textsc{ExtDisrHD}(v, G^*_N, S^*)$ *to be the set of all externally disrespectful nodes in* $\textsc{in}(\text{SC}^{depth}_v)$*. We further define the set* $\textsc{ExtDisrHD}(G^*_N, S^*) =$

$\bigcup_{v \in V(G_N)} \textsc{ExtDisrHD}(v, G_N^*, S^*)$ *of all externally disrespectful nodes in the graph* G_N^*.

We upper bound the number of d'-perfectly connected externally disrespectful nodes in a given superconcentrator as follows.

Lemma 9. *Given a d'-perfectly connected superconcentrator $\textsc{sc}_v^{depth}$ in G_N^* and a set S^* of nodes removed from G_N^*, $|\textsc{ExtDisrHD}(v, G_N^*, S^*)| \leq |S_v^*|$.*

Proof. Suppose, for the sake of contradiction, that $|\textsc{ExtDisrHD}(v, G_N^*, S^*)| > |S_v^*|$. Let $\rho = r(v, G_N^*, S^*)$ be the r value of the given superconcentrator $\textsc{sc}_v^{depth}$. For all $x \in \textsc{ExtDisrHD}(v, G_N^*, S^*)$, let $u_x \in V_n$ be the unique node such that $\texttt{in}_v(u_x) = x$ and let $a_x = out(u_x, v)$. Observe that if such a u_x does not exist, then x cannot be externally disrespectful, by definition. By property (2) of an externally disrespectful node, we know that either (a) $r(u_x, G_N^*, S^*) \geq \rho - 1$ and $|S_{u_x}^*| > 0$ or (b) $r(u_x, G_N^*, S^*) \geq \rho - |d'|$ and $|S_{u_x}^*| = 0$. We separate these nodes $X = \{x \in \textsc{ExtDisrHD}(v, G_N^*, S^*)\}$ into two disjoint sets $X_1 = \{x \in \textsc{ExtDisrHD}(v, G_N^*, S^*) : |S_{u_x}^*| > 0\}$ with (a) $r(u_x, G_N^*, S^*) \geq \rho - 1$ and $X_2 = \{x \in \textsc{ExtDisrHD}(v, G_N^*, S^*) : |S_{u_x}^*| = 0\}$ with (b) $r(u_x, G_N^*, S^*) \geq \rho - d'$, accordingly. We know by property (3) of an externally disrespectful node that none of the nodes a_x are internally disrespectful. Observe that this also means that a_x cannot be in the set S^* of removed nodes since any output node in S^* will be internally disrespectful by definition. Furthermore, by property (1) of an externally disrespectful node the node $x \notin S^*$ i.e., x was not deleted.

This means that $\forall x \in X_1$ we have

$$\begin{aligned} \texttt{depth}(x, G_N^* - _V S^*) &\geq \texttt{depth}(a_x, G_N^* - _V S^*) + 1 && (4)\\ &\geq r(u_x, G_N^*, S^*) + 2 && (5)\\ &\geq \rho + 1\ . && (6) \end{aligned}$$

Where inequality 4 follows because both nodes $x, a_x \notin S^*$ have not been deleted. Inequality 5 follows because a_x is not internally disrespectful, and inequality 6 from property (a), above.

Similarly, $\forall x \in X_2$ we have

$$\begin{aligned} \texttt{depth}(x, G_N^* - _V S^*) &\geq \texttt{depth}(a_x, G_N^* - _V S^*) + 1 && (7)\\ &\geq r(u_x, G_N^*, S^*) + d' + 1 && (8)\\ &\geq \rho + 1\ . && (9) \end{aligned}$$

Where inequality 7 follows because both nodes $x, a_X \notin S^*$ have not been deleted. Inequality 8 follows by Definition 13 because $|S_{u_x}^*| = 0$, and inequality 9 from property (b), above.

Thus, there are at least $|S_v^*| + 1$ nodes x inX with $\texttt{depth}(x, G_N^* - _V S^*) \geq \rho + 1$. It follows from Definition 4 that $r(v, G_n^*, S^*) \geq \rho + 1$. Since we defined $\rho = r(v, G_n^*, S^*)$ this implies that $r(v, G_n^*, S^*) \geq 1 + r(v, G_n^*, S^*)$ a clear contradiction. This contradiction means our original assumption is impossible, and $|\textsc{ExtDisrHD}(v, G_N^*, S^*)| \leq |S_v^*|$. □

Now that we have modified some of the underlying sets of bad nodes for each superconcentrator, we will once again use these sets to define the set S of edges which must be removed from the original graph G_N.

Definition 17. *Given G_N let $G_N^* = \textsc{Transform}(G_N, \mathrm{SC}^{depth}, I, O)$ be the transformed graph constructed as in Definition 3 and let $S^* \in V(G_N^*)$ be any set of nodes. In addition to the set $\textsc{ToExclude}(G_N^*, S^*) \subseteq E(G_N)$ we define three additional sets of edges to remove from G_N (1) $S_{\textsc{Dest}} \doteq \{e \in E(G_N) \ : \ \texttt{in(e)} \in \textsc{Dest}(G_N^*, S^*)\}$, (2) $S_{\textsc{IntDisrHD}} \doteq \{e \in E(G_N) \ : \ \texttt{out(e)} \in \textsc{IntDisrHD}(G_N^*, S^*)\}$, and (3) $S_{\textsc{ExtDisrHD}} \doteq \{e \in E(G_N) \ : \ \texttt{in(e)} \in \textsc{ExtDisrHD}(G_N^*, S^*)\}$. We further define the set of all removed edges $S = \textsc{NtoEHD}(G_N^*, S^*) \doteq \textsc{ToExclude}(G_N^*, S^*) \cup S_{\textsc{Dest}} \cup S_{\textsc{IntDisrHD}} \cup S_{\textsc{ExtDisrHD}}$.*

We can bound the total size of this set S.

Lemma 10. *For any graph G_N^* and any subset $S^* \subseteq V(G_N^*)$ of nodes we have $|\textsc{NtoEHD}(G_N^*, S^*)| \leq 3|S^*|$.*

Proof. Let $\texttt{Spoiled} \subset V(G_N)$ denote the set of all nodes $v \in V(G_N)$ for which $\textsc{Spoiled}(v, G_N^*, S^*) = 1$ i.e., the corresponding super concentrator $\textsc{sc}_v^{depth}$ in G_N^* was 'spoiled' by removing at least $\left|\textsc{in}(\textsc{sc}_v^{depth})\right|$ nodes from $\textsc{sc}_v^{depth}$. We now partition the set S^* into two disjoint subsets, $S_A^* \doteq \bigcup_{v \in \texttt{Spoiled}} S_v^*$ and $S_B^* \doteq \bigcup_{v \notin \texttt{Spoiled}} S_v^* = S^* \setminus S_A^*$ — recall that $S_v^* = \textsc{sc}_v^{depth} \cap S^*$. Let $S = \textsc{NtoE}(G_N^*, S^*)$ be the corresponding set of edges. We can similarly partition S into two sets $S_A = \{(u,v) \in S \ : \ u \in \texttt{Spoiled} \vee v \in \texttt{Spoiled}\}$ and $S_B = \{(u,v) \in S \ : \ u \notin \texttt{Spoiled} \wedge v \notin \texttt{Spoiled}\}$.

We will show that $|S_A| \leq 3\,|S_A^*|$ and $|S_B| \leq 3\,|S_B^*|$. From this it immediately follows that

$$|S| = |S_A| + |S_B| \leq 3\,|S_A^*| + 3\,|S_B^*| \leq 3|S^*| \ .$$

We first note that $S_A = \textsc{ToExclude}(G_N, \textsc{sc}, I, O, S_A^*)$ so it follows immediately from Lemma 1 that $|S_A| \leq 3|S_A^*|$. We now bound $|S_B|$. We first observe that $S_B = (S_{\textsc{Dest}} \cup S_{\textsc{IntDisrHD}} \cup S_{\textsc{ExtDisrHD}}) \setminus \textsc{ToExclude}(G_N, \textsc{sc}, I, O, S_A^*)$. Thus, it suffices to show that $|S_{\textsc{Dest}} \setminus S_A| \leq |S_B^*|$ and $|S_{\textsc{IntDisrHD}} \setminus S_A| \leq |S_B^*|$ and $|S_{\textsc{IntDisrHD}} \setminus S_A| \leq |S_B^*|$.

Given $v \in V(G_N)$ let $S_{\textsc{Dest},v} \doteq \left\{(x,y) \in S_{\textsc{Dest}} \setminus S_A : y \in \textsc{in}(\textsc{sc}_v^{depth})\right\}$ be the set of edges in $S_{\textsc{Dest}}$ which were destroyed because we removed an input node from $\textsc{sc}_v^{depth}$. For every node $v \in V(G_N)$ we have $|S_{\textsc{Dest},v}| \leq |S_v^*| = \gamma_v \left|\textsc{in}(\textsc{sc}_v^{depth})\right|$. Because any edge $(x,y) \in S_{\textsc{Dest}} \setminus S_A$ must have both x and y unspoiled we also have

$$S_{\textsc{Dest}} \setminus S_A = \bigcup_{v \in \overline{\texttt{Spoiled}}} S_{\textsc{Dest},v} \ .$$

It follows immediately that

$$|S_{\textsc{Dest}} \setminus S_A| \leq \sum_{v \in \overline{\texttt{Spoiled}}} |S_{\textsc{Dest},v}| \leq \sum_{v \in \overline{\texttt{Spoiled}}} \gamma_v \left|\textsc{in}(\textsc{sc}_v^{depth})\right| = \sum_{v \in \overline{\texttt{Spoiled}}} |S_v^*| = |S^* \setminus S_A| = |S_B^*|$$

Similarly, given $v \in V(G_N)$ let $S_{\text{IntDisrHD},v} \doteq \left\{(x,y) \in S_{\text{IntDisrHD}} \setminus S_A : x \in \text{OUT}(\text{SC}_v^{depth})\right\}$ be the set of edges in $S_{\text{IntDisrHD}}$ which are internally disrespectful and exit the superconcentrator SC_v^{depth}. For every node $v \in V(G_N)$ we have

$$|S_{\text{IntDisrHD},v}| \leq |\text{IntDisrHD}(v, G_N^*, S^*)| \leq |S_v^*| \leq \gamma_v \left|\text{IN}(\text{SC}_v^{depth})\right| \ .$$

The first inequality follows because each node in $\text{IntDisrHD}(v, G_N^*, S^*) \subseteq \text{OUT}(\text{SC}_v^{depth})$ is incident to exactly one edge from $S_{\text{IntDisrHD},v}$, the second inequality follows by Lemma 8 and the third follows by definition. We now have

$$S_{\text{IntDisrHD}} \setminus S_A = \bigcup_{v \in \overline{\texttt{Spoiled}}} S_{\text{IntDisrHD},v}$$

because any edge $(x,y) \in S_{\text{IntDisrHD}} \setminus S_A$ must have both x and y unspoiled. Thus, it immediately follows that

$$|S_{\text{IntDisrHD}} \setminus S_A| \leq \sum_{v \in \overline{\texttt{Spoiled}}} |S_{\text{IntDisrHD},v}| \leq \sum_{v \in \overline{\texttt{Spoiled}}} \gamma_v \left|\text{IN}(\text{SC}_v^{depth})\right| = \sum_{v \in \overline{\texttt{Spoiled}}} |S_v^*| = |S^* \setminus S_A| = |S_B^*| \ .$$

Finally, given $v \in V(G_N)$ let $S_{\text{ExtDisrHD},v} \doteq \left\{(x,y) \in S_{\text{ExtDisrHD}} \setminus S_A : y \in \text{IN}(\text{SC}_v^{depth})\right\}$ be the set of edges in $S_{\text{ExtDisrHD}}$ which are externally disrespectful and enter the superconcentrator SC_v^{depth}. For every node $v \in V(G_N)$ we have

$$|S_{\text{ExtDisrHD},v}| \leq |\text{ExtDisrHD}(v, G_N^*, S^*)| \leq |S_v^*| \leq \gamma_v \left|\text{IN}(\text{SC}_v^{depth})\right| \ .$$

The first line follows because each node in $\text{ExtDisrHD}(v, G_N^*, S^*) \subseteq \text{IN}(\text{SC}_v^{depth})$ is incident to exactly one edge from $S_{\text{ExtDisrHD}}$, the second inequality follows directly from Lemma 9 and the third follows by definition. We now have

$$S_{\text{ExtDisrHD}} \setminus S_A = \bigcup_{v \in \overline{\texttt{Spoiled}}} S_{\text{ExtDisrHD},v}$$

because any edge $(x,y) \in S_{\text{ExtDisrHD}} \setminus S_A$ must have both x and y unspoiled. Thus, it immediately follows that

$$|S_{\text{ExtDisrHD}} \setminus S_A| \leq \sum_{v \in \overline{\texttt{Spoiled}}} |S_{\text{ExtDisrHD},v}| \leq \sum_{v \in \overline{\texttt{Spoiled}}} \gamma_v \left|\text{IN}(\text{SC}_v^{depth})\right| = \sum_{v \in \overline{\texttt{Spoiled}}} |S_v^*| = |S^* \setminus S_A| = |S_B^*| \ .$$

□

It now remains to show the actual connection between this set S and S^*, intuitively that removing S^* from G_N^* is (nearly) equivalent to removing S from G_N.

Theorem 9. *Given a graph G_N which is (e,d)-*edge *depth robust. Given a family of superconcentrators SC^{depth} with perfect connectivity at depth d', and families $I = \{\texttt{in}_v\}_{v \in V}$ and $O = \{\texttt{out}_v\}_{v \in V}$ of one-to-one functions, let $G_N^* = \text{Transform}(G_N, \text{SC}^{depth}, I, O)$ be the transformed graph constructed as in Definition 3. G_N^* is $(e/3, (d-e) \cdot d')$-*node *depth robust for all $d > 1$.*

Proof. Suppose, for the sake of contradiction, that G^*_N is not $(e/3, (d-e)\cdot d')$-*node* depth robust. Let S^* be the set of at most $e/3$ nodes removed from G^*_N such that $\texttt{depth}(G^*_N - _V S^*) < (d-e)\cdot d')$. Let $S = \textsc{NtoEHD}(G^*_N, S^*)$ be a set of edges in G_N as in Definition 17. We know from Lemma 10 that $|S| \leq 3|S^*|$.

Consider the graph $G_N - _E S$. Since $|S| \leq e$, and G_N is (e,d)-edge depth robust, $\texttt{depth}(G - _E S) \geq d$. Let $P(v_d) = v_1, v_2, \ldots, v_d$ be the path of length d in $G_N - \textsc{NtoE}(G^*_N, S^*)$. By definition, we have $(v_i, v_{i+1}) \in E(G_N - _E S)$ for all $i \in [1, d-1]$. Define $r_i \doteq r(v_i, G^*_N, S^*)$ for each node v_i in this path.

We define the predicate $\textsc{Unaffected}(v, G^*_N, S^*)$ to be true if and only if no node is removed from v's corresponding superconcentrator, i.e. $\textsc{Unaffected}(v, G^*_N, S^*) = (|S^*_v| == 1)$. We will argue that $r_1 \geq 0$, $r_{i+1} \geq r_i + 2$ if $\neg\textsc{Unaffected}(v_i, G^*_N, S^*)$, and $r_{i+1} > r_i + d'$ if $\textsc{Unaffected}(v_i, G^*_N, S^*)$. Since at least $d-e$ nodes in $P(v_d)$ must be unaffected, it will immediately follow that $r_d \geq (d-e-1)d'$ if $\textsc{Unaffected}(v_d, G^*_N, S^*)$ or $r_d \geq (d-e)d'$ if $\neg\textsc{Unaffected}(v_d, G^*_N, S^*)$ — intuitively, if the final node is unaffected, we may miss one node's worth of d' gain.

Since the minimum value of r is 0, we know $r_1 = r(v_1, G^*_N, S^*) \geq 0$. Now let us consider the edge $(v_i, v_{i+1}) \in E(G - _E S)$ for any v_i. We know that $(v_i, v_{i+1}) \notin S_{\textsc{Dest}} \cap S_{\textsc{ExtDisrHD}} \cap S_{\textsc{IntDisrHD}}$ and there must be a corresponding edge $(u^*_i, u^*_{i+1}) \in E(G^*_N)$. Such an edge (v_i, v_{i+1}) would have been deleted if u^*_{i+1} were destroyed, u^*_i were internally disrespectful, u^*_{i+1} were externally disrespectful, or either v_{i+1} or v_i were spoiled.

Therefore, the node $u^*_{i+1} \in \textsc{in}(\textsc{sc}^{depth}_{v_{i+1}})$ must have the following properties: First, $u^*_{i+1} \notin S^*$, i.e. it has not been removed/destroyed. Second, u^*_{i+1}'s parent, u^*_i is not in the set $\textsc{IntDisrHD}(G^*_N, S^*)$ i.e., u^*_i is not internally disrespectful. Third, $u^*_{i+1} \notin \textsc{ExtDisrHD}(G^*_N, S^*)$ i.e., u^*_{i+1} is not externally disrespectful.

Recall that the four properties of an externally disrespectful node u^*_{i+1} in a superconcentrator $\textsc{sc}^{depth} v_{i+1}$, with parent u^*_i in a superconcentrator $\textsc{sc}^{depth}_{v_i}$ are: (1) u^*_{i+1} is not destroyed (2) either (a) $r(v_{i+1}, G^*_N, S^*) \leq r(v_i, G^*_N, S^*) + 1$ with $\neg\textsc{Unaffected}(v_i, G^*_N, S^*)$ or (b) $r(v_{i+1}, G^*_N, S^*) \leq r(v_i, G^*_N, S^*) + d'$ with $\textsc{Unaffected}(v_i, G^*_N, S^*)$ (3) u^*_{i+1}'s parent, u^*_i, is not internally disrespectful (4) neither v_{i+1} nor v_i are spoiled. We know that u^*_{i+1} satisfies the first, third, and fourth properties, so if u^*_{i+1} satisfies the second property then $u^*_{i+1} \in$ would be externally disrespectful. Yet we know this is not the case, since u^*_{i+1} is not externally disrespectful. Thus, we must have $r_{i+1} = r(v_{i+1}, G^*_N, S^*) \geq r(v_i, G^*_N, S^*) + 2 = 2 + r_i$ if $\neg\textsc{Unaffected}(v_i, G^*_N, S^*)$ or $r(v_{i+1}, G^*_N, S^*) \geq r(v_i, G^*_N, S^*) + d' + 1$ if $\textsc{Unaffected}(v_i, G^*_N, S^*)$.

Thus, by induction, we know that $r_d \geq (d-e-1)d'$ if $\textsc{Unaffected}(v_d, G^*_N, S^*)$ or $r_d \geq (d-e)d'$ if $\neg\textsc{Unaffected}(v_d, G^*_N, S^*)$.

Consider the former case that $\textsc{Unaffected}(v_d, G^*_N, S^*)$ and recall that we know $d > 1$. Then there must be some edge (u^*_i, u^*_{i+1}) leading into the final node of $P(v_d)$, which means $\textsc{sc}^{depth}_{v_d}$ is not spoiled (as all edges leading into a spoiled superconcentrator are removed). This, in turn, means that $\gamma_{v_d}(S^*) < \frac{2}{3}$. Now, consider the set of output nodes from $\textsc{sc}^{depth}_{v_d}$, $\textsc{out}(\textsc{sc}^{depth}_{v_d})$. By Definition 13, we know that $\forall x \in \textsc{out}(\textsc{sc}^{depth}_{v_d}), \texttt{depth}(x) \geq r(v_d, G^*_N, S^*) + d' \geq r_d + d'$.

This means that regardless of whether v_d is unaffected, there exists at least one node with depth $\geq (d-e)d'$. This contradiction means our original assumption was impossible, and G^*_N is $(e/3, (d-e)d')$-node depth robust. □

D Proofs of Section 3.4

We wish to construct a concretely efficient graph G^{goal} with N nodes and indegree 2 such that G^{goal} is $(c_1 N/\log N, c_2 N)$-node depth robust for constants c_1, c_2 as large as possible e.g., maximizing the product $c_1 c_2$. To obtain a such a graph by applying our transformation we would need to start with a concretely efficient $(c'_1 N', c'_2 N')$ edge depth-robust graph with N' nodes and $O(N' \log N')$ edges. The Grates graph [25] has $N' = 2^{n+1} - 1$ nodes and $n2^n$ edges and is known to be (e_i, d_i) edge depth robust with $e_i = i \cdot 2^{n-1} - 1$ and $d_i = \frac{2^{n-1}}{\sum_{j=0}^{i} \binom{n}{j}}$ for all $i \leq n$. Unfortunately, this is not sufficient to obtain a $(\Omega(N/\log N), \Omega(N))$-node depth robust graph via our transformation. Even if we plug in $i = 1$ we still have $d_1 = \frac{2^{n-1}}{n+1} = o(N')$ which is already too small for our purposes. In this section we provide a new analysis of the Grates graph and demonstrate that the graph is $(e, 2d-1)$-edge depth robust for any parameters e, d such that $e + d \leq 2^n$, $d > 1$.

We have shown previously shown in Theorem 9 that our transformation can transform any (e, d)-edge depth robust graph into an $(e/3, (d-e)d')$-*node* depth-robust graph for all $e \leq \frac{1}{2}d$, $d > 1$. However, for any graph G_N, our transformation does increase the number of nodes in $G^*_N = \textsc{Transform}(G_N, \textsc{sc}, I, O)$, such that $|V(G^*_N)| > |E(G_N)|$. So we could construct such a G^{goal} graph with an underlying edge-depth-robust graph G^{base} if G^{base} has the following properties:

(1) G^{base} has $\Omega(N')$ edges, (2) G^{base} is $(\Omega(N'), \Omega(N'/\log N'))$-edge depth-robust. (3) G^{base} has indegree $\log N$.

To do so, we will use the Grates graph, first introduced by [25]. This is a simple, well-studied, edge-depth robust graph. The Grates graph is defined as follows:

Definition 18. *We define the graph* $G^{\texttt{Grates}}_n = (V(G^{\texttt{Grates}}_n), E(G^{\texttt{Grates}}_n))$*, as in [25], as follows:*
$V(G^{\texttt{Grates}}_n) \doteq \bigcup_{i=0}^{n} \{0,1\}^i$
$E(G^{\texttt{Grates}}_n) \doteq \{(u0w, u) \mid |u0w| = n\} \cup \{(u, u1w) \mid |u1w| = n\}$.

Figure 5 depicts the Grates graph with $n = 3$:

Intuitively, we can view each node as part of a binary tree, although the Grates graph does not include the same edges. Instead, there is a directed edge connecting a leaf node with all of its ancestors. The direction is determined by whether the leaf appears to the left or the right.

We formalize this concept of the ancestors of a node and the least common ancestors of a pair of nodes in Definition 19. In Figure 1, nodes 100 and 110 are depicted in blue, while $\mathbf{LCA}(100, 110) - 1 -$ is shown in red.

Definition 19. *Given* $u, v \in \{0,1\}^{\leq n} = V(G^{\texttt{Grates}}_n)$*. We say that* u *is an ancestor of* v *if there exists a string* $w \in \{0,1\}^*$ *such that* $v = u \| w$ *(here* $\|$ *denotes*

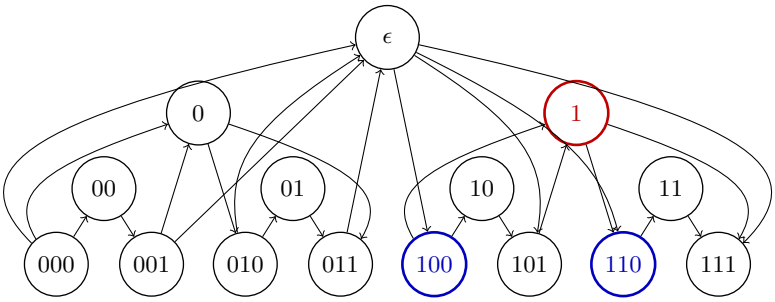


Fig. 5: Figure 5: $G_n^{\texttt{Grates}}$ with $n = 3$

string concatenation). Given two nodes $x, y \in \bigcup_{i=1}^{n} \{0,1\}^i$ *we define the* least common ancestor $\mathbf{LCA}(x, y)$ *to be the longest string* $z \in \{0,1\}^{\leq n}$ *such that* z *is an ancestor of* x *and* y.

We make a few observations about $G_n^{\texttt{Grates}}$: (1) $|V(G_n^{\texttt{Grates}})| = 2^{n+1} - 1$ (2) $|E(G_n^{\texttt{Grates}})| = n2^n$ (3) $\texttt{indeg}(G_n^{\texttt{Grates}}) = 2^{n-1}$.

First, we give a simple proof that Grates is edge-depth robust. We note that while [25] gives a definition showing $G_n^{\texttt{Grates}}$ is $\left(i \cdot 2^{n-1}, \frac{|V_n|}{4(n+1)}\right)$ for all $i \leq n$, our proof instead shows that $G_n^{\texttt{Grates}}$ is (e, d)-edge depth-robust for any e, d such that $e + d \leq 2^n$, improving the bounds in the case $i = 1$. In [24], Schnitger also transforms $G_n^{\texttt{Grates}}$ — for any particular ϵ — into a (e, d)-*node* depth-robust graph $H_{n,\epsilon}$ with $e = \Omega(n2^n)$, $d = 2^{(1-\epsilon)n}$. However, the transformation is not concretely efficient and depends on the particular constant ϵ. In particular, as $\epsilon \to 0$ the transformation becomes less and less efficient.

Reminder of Theorem 2. $G_n^{\texttt{Grates}}$ *(defined in Definition 18) is* $(e, 2d-1)$*-edge depth-robust for all depth-robustness parameters* e, d *such that* $e + d \leq 2^n$.

Proof. Let S be any set of edges in $G_n^{\texttt{Grates}}$ of size $|S| = e$. Consider the set of all leaf nodes in $G_n^{\texttt{Grates}}$, i.e. $\texttt{Leaf} = \{0,1\}^n$. Intuitively, we consider leaf node $v \in \texttt{Leaf}$ to *destroyed* if it is incident to an edge $e \in S$ that is being removed. We define $\texttt{Leaf}_{\text{DEST}} = \{w \in \{0,1\}^n : \exists y \in \bigcup_{i=0}^{n} \{0,1\}^n . ((y, w) \in S \vee (w, y) \in S)\}$ to be the set of all *destroyed* leaf nodes and we define $\texttt{Leaf}_{\text{GOOD}} = \texttt{Leaf} \setminus \texttt{Leaf}_{\text{DEST}}$ to be all of the remaining leaf nodes that were not *destroyed*.

Now consider any pair $u, v \in \texttt{Leaf}_{\text{GOOD}}$ of leaf nodes that have not been destroyed and (wlog) assume that u comes before v in a topological ordering of $G_n^{\texttt{Grates}}$. Observe that the path $u \to \mathbf{LCA}(u, v) \to v$ is still a valid path in $G_n^{\texttt{Grates}} - _E S$. In particular, we have $(u, \mathbf{LCA}(u, v)) \in E(G_n^{\texttt{Grates}}) \setminus S$ because if $(u, \mathbf{LCA}(u, v)) \in S$ then, by definition, we would not have $u \in \texttt{Leaf}_{\text{GOOD}}$. Similarly, $(\mathbf{LCA}(u, v), v) \in E(G_n^{\texttt{Grates}}) \setminus S$.

We know that $|\texttt{Leaf}_{\text{GOOD}}| \geq 2^n - |S|$ because each edge in S can be incident to at most one leaf node — by construction $G_n^{\texttt{Grates}}$ does not contain any edge directly connecting leaf nodes. Let $v_1, \ldots, v_{2^n - |S|}$ be distinct nodes in $\texttt{Leaf}_{\text{GOOD}}$ listed in topological order. For each $1 \leq i < 2^n - |S|$ we have a direct path $v_i \to$

$\mathbf{LCA}(v_i, v_{i+1}) \to v_{i+1}$ in $G_n^{\texttt{Grates}} - _E S$. Thus, $G_n^{\texttt{Grates}} - _E S$ has a path containing *at least* $2^n - |S| + 2^n - |S| - 1$ nodes. □

We *can* show that $\textsc{Transform}(G_n^{\texttt{Grates}}, \text{SC}^{depth}, I, O)$ is node-depth robust by applying Theorem 9 and 2. In particular, applying Theorem 9 we can conclude that $\textsc{Transform}(G_n^{\texttt{Grates}}, \text{SC}^{depth}, I, O)$ is $(e/3, (d-e)d')$-depth robust for any $e + d \leq 2^n$. For our analysis, we will select $d' = n = \log N$ so that size of the transformed graph $\textsc{Transform}(G_n^{\texttt{Grates}}, \text{SC}^{depth}, I, O)$ has $O(N \log N)$ nodes.

Theorem 10. *Let $G_n^{\texttt{Grates}}$ be the Grates graph. Given a family of high-depth superconcentrators SC^{depth} for the depth $d' = n$, and families $I = \{\texttt{in}_v\}_{v \in V}$ and $O = \{\texttt{out}_v\}_{v \in V}$ of one-to-one functions, let $G_N^* = \textsc{Transform}(G_n^{\texttt{Grates}}, \text{SC}^{depth}, I, O)$ be the transformed graph constructed as in Definition 3. G_N^* is $(e/3, 2nd - ne - n)$-*node *depth robust for all $e + d \leq 2^n$, $d > 1$.*

Proof. By Theorem 2, $G_n^{\texttt{Grates}}$ is $(e, 2d-1)$-*edge* depth-robust for $e + d \leq 2^n$. By applying Theorem 9, we therefore know $\textsc{Transform}(G_n^{\texttt{Grates}}, \text{SC}^{depth}, I, O)$ is $(e/3, ((2d-1) - e)\, d')$-*node* depth-robust. This simplifies to $(e/3, 2nd - ne - n)$. □

However, such a transformed Grates graph, like previous transformed graphs, increases in size. It remains to analyze this increase to show that our gain in depth-robustness now keeps pace with our gain in node count.

Lemma 11. *Given a depth n, family of superconcentrators SC^{depth} with perfect connectivity at depth d' and families $I = \{\texttt{in}_v\}_{v \in V}$ and $O = \{\texttt{out}_v\}_{v \in V}$ of one-to-one functions, let $G_N^* = \textsc{Transform}(G_n^{\texttt{Grates}}, \text{SC}^{depth}, I, O)$ be the transformed graph constructed as in Definition 3. G_N^* has $\Sigma_{i=0}^{n-1} 2^i |V(\text{SC}^{depth}(2^{n-i-1})| + \Sigma_{L=0}^{n} \binom{n}{L} |V(\text{SC}^{depth}(max(L, n-L))|$ nodes.*

Proof. The initial Grates graph has $|V(G_n^{\texttt{Grates}})| = 2^{n-1} - 1$. Each of these nodes $v \in V(G_n^{\texttt{Grates}})$ is replaced with a d'-perfectly connected $\texttt{mdeg}(v)$-superconcentrator. Intuitively, we calculate the total number of nodes by counting how many nodes the superconcentrator corresponding to each vertex at each level of the Grates graph contributes. To do so, we first analyze the `mdeg` of each node, since this controls the size of superconcentrator used. For any internal (non-leaf) node v at height i, this is half the number of descendants — since each descendant is connected, and these edges are split evenly between in edges and out edges — i.e. $\texttt{mdeg}(v) = 2^{n-i-1}$.

For a particular leaf node v, finding $\texttt{mdeg}(v)$ is slightly more complicated; while a naive upper bound gives us $\texttt{mdeg}(v) \leq n$, since only one edge connects a particular leaf to each level, this bounds the *total* degree, not `mdeg`. Let L be the number of incoming (leftward) edges, and $R \doteq n - L$ be the number of outgoing (rightward) edges. Then $\texttt{mdeg}(v) = \max(L, n - L)$. For a particular value of L there are $\binom{n}{L}$ possible nodes with that many incoming edges, so we can upper bound the total `mdeg` by summing over this value.

Formally, the total number of nodes in $\textsc{Transform}(G_n^{\texttt{Grates}}, \textsc{sc}^{depth}, I, O)$ is given by:

$$|V(G_N^*)| = \Sigma_{v \in V(G_n^{\texttt{Grates}})} |V(\textsc{sc}_v^{depth})| \tag{10}$$

$$= \Sigma_{i=0}^{n-1} \Sigma_{v \in \{0,1\}^i} |V(\textsc{sc}_v^{depth})| + \Sigma_{v \in \{0,1\}^n} |V(\textsc{sc}_v^{depth})| \tag{11}$$

$$= \Sigma_{i=0}^{n-1} 2^i |V(\textsc{sc}^{depth}(2^{n-i-1})| + \tag{12}$$

$$\Sigma_{L=0}^{n} \binom{n}{L} |V(\textsc{sc}^{depth}(max(L, n-L))| \tag{13}$$

$$\tag{14}$$

□

Lemma 11 still abstracts away the size of a given Superconcentrator. Linear-sized superconcentrators have been shown to exist, including randomized constructions of superconcentrators with $|E(\textsc{sc}(n))|/n)| = 25.3$ [19] and explicit constructions of superconcentrators with $|E(\textsc{sc}(n))|/n = 44$ [1]. By applying indegree reduction to them, these may be transformed to superconcentrators with the following properties:

Lemma 12. *Suppose that* $\textsc{sc}'(N)$ *is a* N*-superconcentrator. Then there is a* N*-superconcentrator* $\textsc{sc}(N)$ *with* $|V(\textsc{sc}(N))| \leq N + |E(\textsc{sc}'(N))|$ *nodes such that* $\textsc{sc}(N)$ *has maximum indegree* 2*. Moreover, the transformation from* $\textsc{sc}'(N)$ *to* $\textsc{sc}(N)$ *is explicit so if* $\textsc{sc}'(N)$ *is explicitly constructible then so is* $\textsc{sc}(N)$*.*

Proof. We can apply indegree reduction and replace each node $v \in V(\textsc{sc}^{dense\prime}(N))$ with a path of length $\texttt{indeg}(v)$ and then distribute incoming edges to v across the path. This ensures that the indegree reduced graph $\textsc{sc}^{dense}(N)$ has maximum indegree 2. The reduction adds

$$\sum_{v:\texttt{indeg}(v)\geq 1} (\texttt{indeg}(v)-1) = |E(\textsc{sc}^{dense\prime}(N))| - \left|\left\{v \in V(\textsc{sc}^{dense\prime}(N)) : \texttt{indeg}(v) \geq 1\right\}\right| ,$$

new nodes to the graph. We can assume WLOG that the only nodes with $\texttt{indeg}(v) = 0$ are the N input nodes — if $\texttt{indeg}(v) = 0$ and $v \in V(\textsc{sc}^{dense\prime}(N))$ is not an input it can be deleted without loss of generality. Thus, the total number of nodes in the graph is at most $|V(\textsc{sc}(N))| \leq N + |E(\textsc{sc}'(N))|$. □

Lemma 13. *Suppose that* $\textsc{sc}'(N)$ *is a* N*-superconcentrator. Then for all* $d' \geq 1$ *there is an explicit transformation from* $\textsc{sc}'(N)$ *to* $\textsc{sc}(N)$ *such that (1)* $\textsc{sc}(N)$ *is an* N*-superconcentrator, (2)* $\textsc{sc}(N)$ *is* d'*-perfectly connected (3)* $\textsc{sc}(N)$ *has indegree* 2*, (4) and* $|V(\textsc{sc}(N))| \leq |V(\textsc{sc}'(N))| + 2N + d'$*.*

Proof. We first apply Lemma 12 to transform the graph $\textsc{sc}'(N)$ to a graph $\textsc{sc}''(N)$ such that (1) $\textsc{sc}''(N)$ is still an N-superconcentrator. (2) $\textsc{sc}''(N)$ has maximum indegree 2 and (3) $|V(\textsc{sc}''(N))| \leq |V(\textsc{sc}'(N))| + N$. We can then apply Claim 8 to transform $\textsc{sc}''(N)$ into a new graph $\textsc{sc}(N)$ by adding $N + d'$ new nodes. By Claim 8, the resulting graph $\textsc{sc}(N)$ is both an N-superconcentrator and is d'-perfectly connected. The final graph has $|V(\textsc{sc}(N))| \leq |V(\textsc{sc}(N))| + 2N + d'$ total nodes and maximum indegree 2. □

Fact 11 *[19] There exists a family of superconcentrators* SC^{dense} *such that, for all* N*: (1)* $\textsc{sc}^{dense}(N)$ *is an* N*-superconcentrator, and (2)* $|E(\textsc{sc}^{dense}(N))| \leq 25.3N$.

Fact 12 *[1] There exists a family of superconcentrators* $\mathrm{SC}^{explicit}$ *such that, for all* N*: (1)* $\textsc{sc}^{explicit}(N)$ *is an* N*-superconcentrator, (2)* $|E(\textsc{sc}^{explicit}(N))| \leq 44N$*, and (3)* $\textsc{sc}^{explicit}(N)$ *is explicitly defined.*

Next, we transform these families to be d'-perfectly connected with maximum indegree 2. Note that since our transformation from Definition 14 retains `indeg` at most 2, these two steps could be taken in either order.

Lemma 14. *Given a depth parameter* d'*, there exists a family of* d'*-perfectly connected superconcentrators* $\mathrm{SC}^{dense\&depth}$ *such that, for all* N*: (1)* $\textsc{sc}^{dense\&depth}(N)$ *is an* N*-superconcentrator, (2)* $\textsc{sc}^{dense\&depth}(N)$ *is* d'*-perfectly connected (3)* $\textsc{sc}^{dense\&depth}(N)$ *has indegree* 2*, (4) and* $|V(\textsc{sc}^{dense\&depth}(N))| \leq 27.3N + d'$.

Proof. The claim follows immediately by applying Lemma 13 to the superconcentrator family from Fact 11. □

Lemma 15. *Given a depth parameter* d'*, there exists a family of* d'*-perfectly connected superconcentrators* $\mathrm{SC}^{explicit\&depth}$ *such that, for all* N*: (1)* $\textsc{sc}^{explicit\&depth}(N)$ *is an* N*-superconcentrator, (2)* $\textsc{sc}^{explicit\&depth}(N)$ *is* d'*-perfectly connected (3)* $\textsc{sc}^{explicit\&depth}(N)$ *has indegree* 2*, (4) and* $|V(\textsc{sc}^{explicit\&depth}(N))| \leq 46N + d'$.

Proof. The claim follows immediately by applying Lemma 13 to the superconcentrator family from Fact 12. □

It remains only to apply these superconcentrators to Theorem 10. First, we give a general form for any linear-size d'-perfectly connected superconcentrator.

Theorem 13. *Given a height* n*, and families* $I = \{\mathtt{in}_v\}_{v \in V}$ *and* $O = \{\mathtt{out}_v\}_{v \in V}$ *of one-to-one functions, let* $d' = n$*. Let* $\textsc{sc}^{depth}$ *be any family of linear-size, indegree* 2*,* d'*-perfectly connected superconcentrators such that, for all* N*,* $|V(\textsc{sc}^{depth}(N)| \leq cN + d'$*. Let* $G^*_N = \textsc{Transform}(G_n^{\mathtt{Grates}}, \textsc{sc}^{depth}, I, O)$ *be the transformed graph constructed as in Definition 3. Then* G^*_N *is an indegree* 2 *graph with* $|V(G^*_N)| \leq (c+2)n2^n + \lceil \frac{n}{2} \rceil \binom{n}{\lfloor n/2 \rfloor} - n$ *which is* $(\frac{e}{3}, (2 \cdot 2^n - 3e - 1)n)$*-node depth robust for all* $e \leq 2^n$*,* $d > 1$.

Proof. By Definition 3, we maintain indegree at most 2. By Theorem 10, G^*_N is $(e/3, (2d - 1 - e)\, n)$*-node* depth-robust for all $e + d \leq 2^n$, $d > 1$. Setting $d = 2^n - e$, we have a $(\frac{e}{3}, (2 \cdot 2^n - 3e - 1)n)$-node depth robust graph.

By Lemma 11, $|V(G^*_N)| \leq \Sigma_{i=0}^{n-1} 2^i |V(\text{SC}^{depth}(2^{n-i-1})| + \Sigma_{L=0}^{n} \binom{n}{L} |V(\text{SC}^{depth}(\max(L, n-L))|$. Using Lemma 16, this simplifies to:

$$
\begin{aligned}
|V(G^*_N)| &= \Sigma_{i=0}^{n-1} 2^i |V(\text{SC}^{depth}(2^{n-i-1})| + && (15)\\
&\quad \Sigma_{L=0}^{n} \binom{n}{L} |V(\text{SC}^{depth}(\max(L, n-L))| && (16)\\
&= \Sigma_{i=0}^{n-1} 2^i \left(c \cdot (2^{n-i-1}) + d'\right) + && (17)\\
&\quad \Sigma_{L=0}^{n} \binom{n}{L} \left(c \cdot \max(L, n-L) + d'\right) && (18)\\
&= \Sigma_{i=0}^{n-1} 2^i d' + \Sigma_{i=0}^{n-1} 2^i c \cdot (2^{n-i-1}) + && (19)\\
&\quad \Sigma_{L=0}^{n} \binom{n}{L} \left(c \cdot \max(L, n-L)\right) + d' 2^n && (20)\\
&= (2^n - 1) d' + \Sigma_{i=0}^{n-1} c \cdot (2^{n-1}) + && (21)\\
&\quad cn2^{n-1} + \lceil \frac{n}{2} \rceil \binom{n}{\lfloor n/2 \rfloor} + d' 2^n && (22)\\
&= d' 2^n - d' + \frac{c}{2} n 2^n + && (23)\\
&\quad \frac{c}{2} n 2^n + \lceil \frac{n}{2} \rceil \binom{n}{\lfloor n/2 \rfloor} + d' 2^n && (24)\\
&= cn2^n + 2d' 2^n + \lceil \frac{n}{2} \rceil \binom{n}{\lfloor n/2 \rfloor} - d' && (25)\\
&= (c+2) n 2^n + \lceil \frac{n}{2} \rceil \binom{n}{\lfloor n/2 \rfloor} - n && (26)
\end{aligned}
$$

□

Corollary 1. *Given a height n, and families $I = \{\texttt{in}_v\}_{v \in V}$ and $O = \{\texttt{out}_v\}_{v \in V}$ of one-to-one functions, let $d' = n$. Let $\text{SC}^{dense\&depth}$ be the family of d'-perfectly connected superconcentrators given by Lemma 14 . Let $G^*_N = \textsc{Transform}(G_n^{\texttt{Grates}}, \text{SC}^{dense\&depth}, I, O)$ be the transformed graph constructed as in Definition 3. Then G^*_N is an indegree 2 graph with $|V(G^*_N)| \leq (29.3)n2^n + \lceil \frac{n}{2} \rceil \binom{n}{\lfloor n/2 \rfloor} - n$ which is $(\frac{e}{3}, (2 \cdot 2^n - 3e - 1)n)$-*node *depth robust for all $e \leq 2^n$, $d > 1$.*

Proof. By Lemma 14, $\text{SC}^{dense\&depth}$ is a family of indegree 2, d'-perfectly connected superconcentrators such that, for all N, $|V(\text{SC}^{depth}(N)| \leq 27.3N + d'$, so the proof follows immediately from Theorem 13. □

Setting $N' = n2^n$, we observe that G^*_N has $c_1 N' + o(N')$ nodes with $c_1 = 29.3$. For the sake of simplicity we will drop lower order terms in the discussion below i.e., assume G^*_N has exactly N' nodes and that is approximately $(e/3, (2 \cdot 2^n - 3e)n)$-node depth-robust for any $e < 2^n$. Recall that the product $(e/3) \cdot (2 \cdot 2^n - 3e)n$ is a lower bound on the cumulative pebbling complexity of G^*_N [4]. The product is maximized by setting $e = \frac{2^n}{3}$, yielding a lower bound of $\frac{(2^n)^2 n}{9}$.

Relating this term back to the number of nodes $c_1 N'$ in the graph G^*_N the lower bound becomes

$$\frac{(2^n)^2\, n}{9} = \frac{N'^2}{9n} = \frac{(c_1 N')^2}{9c_1^2 n} \geq \frac{(c_1 N')^2}{9c_1^2 \log N'} \ .$$

We have $\frac{1}{9c_1^2} \approx 1.29 \times 10^{-4}$. By comparison, the best previously known constant for graphs satisfying

$$ed = \Omega\left(\frac{|V(G)|^2}{\log |V(G)|}\right)$$

was 7.3×10^{-6} for DRSample [3, Thm. 3.1]. Thus, our construction improves the concrete constant by roughly a factor of 18.

Similarly, we can use the superconcentrator from Lemma 15 in Theorem 13.

Corollary 2. *Given a height n, and families $I = \{\mathtt{in}_v\}_{v \in V}$ and $O = \{\mathtt{out}_v\}_{v \in V}$ of one-to-one functions, let $d' = n$. Let $\mathrm{SC}^{explicit\&depth}$ be the family of d'-perfectly connected superconcentrators given by Lemma 15 . Let $G^*_N = \textsc{Transform}(G_n^{\mathtt{Grates}}, \mathrm{SC}^{explicit\&depth}, I, O)$ be the transformed graph constructed as in Definition 3. Then G^*_N is an indegree 2 graph with $|V(G^*_N)| \leq (48)n2^n + \lceil \frac{n}{2} \rceil \binom{n}{\lfloor n/2 \rfloor} - n$ which is $(\frac{e}{3}, (2 \cdot 2^n - 3e - 1)n)$-node depth robust for all $e \leq 2^n$, $d > 1$.*

Proof. By Lemma 14, $\mathrm{SC}^{dense\&depth}$ is a family of indegree 2, d'-perfectly connected superconcentrators such that, for all N, $|V(\mathrm{SC}^{depth}(N)| \leq 46N + d'$, so the proof follows immediately from Theorem 13 □

Once again, we can set $N' = n2^n$ and observe that (dropping lower-order terms) G^*_N has (approximately) $c_1 N'$ nodes with $c_1 = 48$, and is (approximately) $(e/3, (2 \cdot 2^n - 3e)n)$-node depth-robust for any $e < 2^n$. As before we look to lower bound the product $(e/3)((2 \cdot 2^n - 3e)n)$ since this is a lower bound for the cumulative pebblign complexity of G^*_N. The product is maximized at $e = 2^n/3$ so we get a lower bound of $\frac{2^{2n} n}{9}$ as before.

Relating this term back to the number of nodes $c_1 N'$ in the graph G^*_N the lower bound becomes

$$\frac{(2^n)^2\, n}{9} = \frac{N'^2}{9n} = \frac{(c_1 N')^2}{9c_1^2 n} \geq \frac{(c_1 N')^2}{9c_1^2 \log N'} \ .$$

For explicit constructions with $c_1 = 48$ we now have $\frac{1}{9c_1^2} \approx 4.82 \times 10^{-5}$. By comparison, the DRSample constant is 7.3×10^{-6} [3, Thm. 3.1]. Thus, even for the explicit construction, this improves the concrete constant by roughly a factor of 6.6.

D.1 Useful Binomial Identity

The following binomial identity is useful in our analysis of the number of edges in the grates graph. We used chatGPT to find/prove the identify. The authors have verified the proof.

Lemma 16. *For every integer $n \geq 0$,*

$$\sum_{L=0}^{n} \binom{n}{L} \max(L, n-L) = n2^{n-1} + \left\lceil \frac{n}{2} \right\rceil \binom{n}{\lfloor n/2 \rfloor}.$$

Proof. We begin with the identity

$$\max(L, n-L) = \frac{n}{2} + \left| L - \frac{n}{2} \right|.$$

Thus,

$$\sum_{L=0}^{n} \binom{n}{L} \max(L, n-L) = \frac{n}{2} \sum_{L=0}^{n} \binom{n}{L} + \sum_{L=0}^{n} \binom{n}{L} \left| L - \frac{n}{2} \right|.$$

Using

$$\sum_{L=0}^{n} \binom{n}{L} = 2^n,$$

the first term becomes

$$n2^{n-1}.$$

Hence,

$$\sum_{L=0}^{n} \binom{n}{L} \max(L, n-L) = n2^{n-1} + \sum_{L=0}^{n} \binom{n}{L} \left| L - \frac{n}{2} \right|.$$

It remains to evaluate the second sum.

If $n = 2m$ is even, then by symmetry,

$$\sum_{L=0}^{2m} \binom{2m}{L} |L - m| = 2 \sum_{L=m+1}^{2m} \binom{2m}{L} (L - m).$$

Using the identity

$$L\binom{2m}{L} = 2m\binom{2m-1}{L-1},$$

we obtain

$$\sum_{L=m+1}^{2m} \binom{2m}{L} (L - m) = 2m \sum_{L=m+1}^{2m} \binom{2m-1}{L-1} - m \sum_{L=m+1}^{2m} \binom{2m}{L}.$$

Since

$$\sum_{L=m+1}^{2m} \binom{2m-1}{L-1} = 2^{2m-2},$$

and

$$\sum_{L=m+1}^{2m} \binom{2m}{L} = 2^{2m-1} - \frac{1}{2}\binom{2m}{m},$$

it follows that

$$\sum_{L=m+1}^{2m} \binom{2m}{L}(L-m) = \frac{m}{2}\binom{2m}{m}.$$

Therefore,

$$\sum_{L=0}^{2m} \binom{2m}{L}|L-m| = m\binom{2m}{m}.$$

Combining terms gives

$$\sum_{L=0}^{2m} \binom{2m}{L}\max(L, 2m-L) = m2^{2m} + m\binom{2m}{m}.$$

Now suppose that $n = 2m+1$ is odd. Then

$$\left|L - \frac{2m+1}{2}\right| = \frac{1}{2}|2L-(2m+1)|.$$

By symmetry,

$$\sum_{L=0}^{2m+1} \binom{2m+1}{L}\left|L - \frac{2m+1}{2}\right| = 2\sum_{L=m+1}^{2m+1} \binom{2m+1}{L}\left(L - \frac{2m+1}{2}\right).$$

Proceeding as above,

$$\sum_{L=0}^{2m+1} \binom{2m+1}{L}\left|L - \frac{2m+1}{2}\right| = (m+1)\binom{2m+1}{m}.$$

Therefore,

$$\sum_{L=0}^{2m+1} \binom{2m+1}{L}\max(L, 2m+1-L) = (2m+1)2^{2m} + (m+1)\binom{2m+1}{m}.$$

Both cases can be written compactly as

$$\sum_{L=0}^{n} \binom{n}{L}\max(L, n-L) = n2^{n-1} + \left\lceil \frac{n}{2} \right\rceil \binom{n}{\lfloor n/2 \rfloor}.$$

□