A New Connection
Between Node and Edge Depth Robust Graphs

November 6, 2020

Abstract

Given a directed acyclic graph (DAG) $G=(V,E)$, we say that G is (e,d)-depth-robust (resp. (e,d)-edge-depth-robust) if for any set $S\subseteq V$ (resp. $S\subseteq E$) of at most $|S|\leq e$ nodes (resp. edges) the graph $G-S$ contains a directed path of length d. While edge-depth-robust graphs are potentially easier to construct, many applications in cryptography require node depth-robust graphs with small indegree. We create a graph reduction that transforms an (e,d)-edge-depth-robust graph with m edges into a $(e/2,d)$-depth-robust graph with $O(m)$ nodes and constant indegree. One immediate consequence of this result is the first construction of a provably $(\frac{n\log\log n}{\log n},\frac{n}{\log n(\log n)^{\log\log n}})$-depth-robust graph with constant indegree. Our reduction crucially relies on ST-robust graphs, a new graph property we introduce which may be of independent interest. We say that a directed, acyclic graph with n inputs and n outputs is (k_1,k_2)-ST-robust if we can remove any k_1 nodes and there exists a subgraph containing at least k_2 inputs and k_2 outputs such that each of the k_2 inputs is connected to all of the k_2 outputs. If the graph if $(k_1,n-k_1)$-ST-robust for all $k_1\leq n$ we say that the graph is maximally ST-robust. We show how to construct maximally ST-robust graphs with constant indegree and $O(n)$ nodes. Given a family $\mathbb{M}$ of ST-robust graphs and an arbitrary (e,d)-edge-depth-robust graph G we construct a new constant-indegree graph Reduce$(G,\mathbb{M})$ by replacing each node in G with an ST-robust graph from $\mathbb{M}$. We also show that ST-robust graphs can be used to construct (tight) proofs-of-space and (asymptotically) improved wide-block labeling functions.

1 Introduction

Given a directed acyclic graph (DAG) $G=(V,E)$, we say that G is (e,d)-reducible (resp. (e,d)-edge reducible) if there is a subset $S\subseteq V$ (resp. $S\subseteq E$) of $|S|\leq e$ nodes (resp. edges) such that $G-S$ does not contain a directed path of length d. If a graph is not (e,d)-reducible (resp. (e,d)-edge reducible) we say that the graph is (e,d)-depth robust (resp. (e,d)-edge-depth-robust). Depth robust graphs have found many applications in the field of cryptography in the construction of proofs of sequential work [MMV13], proofs of space [DFKP15, Pie19], and in the construction of data independent memory hard functions (iMHFs). For example, highly depth robust graphs are known to be necessary [AB16] and sufficient [ABP17] to construct iMHFs with high amortized space time complexity. While edge depth-robust graphs are often easier to construct [Sch83], most applications require node depth-robust graphs with small indegree.

It has been shown [Val77] that in any DAG with m edges and n nodes, there exists a set S_i of $\frac{mi}{\log n}$ edges that will force $\mathrm{depth}(G-S_i) \le \frac{n}{2^i}$ for all $i < \log n$. For DAGs with constant indegree we have $m = O(n)$ edges so an equivalent condition holds for node depth robustness [AB16], since a node can be removed by removing all the edges incident to it. In particular, there exists a set S_i of $O(\frac{ni}{\log n})$ nodes such that $\mathrm{depth}(G-S_i) \le \frac{n}{2^i}$ for all $i < \log n$. It is known how to construct a $(c_1 n/\log n, c_2 n)$-depth-robust graph, for suitable $c_1, c_2 > 0$ [ABP17] and a $(c_3 n, c_4 n^{1-\epsilon})$-depth-robust graph for small ϵ for [Sch83].

An open challenge is to construct constant indegree $(c_1 ni/\log n, c_2 n/2^i)$-depth-robust graphs which match the Valiant bound [Val77] for intermediate values of $i = \omega(1)$ and $i = o(\log n)$. For example, when $i = \log\log n$ then the Valiant bound [Val77] does not rule out the existence of $(c_1 ni/\log n, c_2 n/\log n)$-depth-robust graphs with constant indegree. Such a graph would yield asymptotically stronger iMHFs [BHK+19]. While there are several constructions that are conjectured to be $(c_1 ni/\log n, c_2 n/\log n)$-depth-robust the best provable lower bound for $(e = cni/\log n, d)$-depth robustness of a constant indegree graph is $d = \Omega(n^{1-\epsilon})$. For edge-depth robustness we have constructions of graphs with $m = O(n \log n)$ edges which are (e_i, d_i)-edge depth robust for any i with $e_i = mi/\log n$ and $d_i = n/\log^{i+1} n$ — much closer to matching the Valiant bound [Val77].

1.1 Contributions

Our main contribution is a graph reduction that transforms any (e,d)-edge-depth-robust graph with m edges into an $(e/2,d)$-depth-robust graph with $O(m)$ nodes and constant indegree. Our reduction utilizes ST-robust graphs, a new graph property we introduce and construct. We believe that ST-robust graphs may be of independent interest.

Intuitively, a (k_1,k_2)-ST-robust graph with n inputs I and n outputs O satisfies the property that, even after deleting k_1 nodes from the graph we can find k_2 inputs $x_1,\dots,x_{k_2}$ and k_2 outputs $y_1,\dots,y_{k_2}$ such that *every* input x_i ($i \in [k_2]$) is still connected to *every* output y_j ($j \in [k_2]$). If we can guarantee that the each directed path from x_i to y_j has length d then we say that the graph is (k_1,k_2,d)-ST-Robust.

A maximally depth-robust graph should be $(k_1, n-k_1)$ -depth robust for any k_1.

Definition 1.1. ST-Robust Let $G=(V,E)$ be a DAG with n inputs, denoted by set I and n outputs, denoted by set O. Then G is (k_1,k_2)-ST-robust if $\forall D \subset V(G)$ with $|D| \le k_1$, there exists subgraph H of $G-D$ with $|I \cap V(H)| \ge k_2$ and $|O \cap V(H)| \ge k_2$ such that $\forall s \in I \cap V(H)$ and $\forall t \in O \cap V(H)$ there exists a path from s to t in H. If $\forall s \in I \cap V(H)$ and $\forall t \in O \cap V(H)$ there exists a path from s to t of length $\ge d$ then we say that G is (k_1,k_2,d)-ST-robust.

Definition 1.2. Maximally ST-Robust Let $G=(V,E)$ be a constant indegree DAG with n inputs and n outputs. Then G is c_1-maximally ST-robust (resp. c_1 max ST-robust with depth d) if there exists a constant $0 < c_1 \le 1$ such that G is $(k,n-k)$-ST-robust (resp. $(k,n-k,d)$-ST-robust) for all k with $0 \le k \le c_1 n$. If $c_1 = 1$, we just say that G is maximally ST-robust.

We show how to construct maximally ST-robust graphs with constant indegree and $O(n)$ nodes and we show how maximally ST-robust graphs can be used to transform any (e,d)-edge-depth-robust graph G with m edges into a $(e/2,d)$-depth-robust graph G' with $O(m)$ nodes and constant indegree. Intuitively, in our reduction each node $v \in V(G)$ with degree δ is replaced with a maximally ST-robust graph M_v with δ inputs/outputs. Incoming edges into v are redirected into the inputs I_v of the ST-robust graph. Similarly, v's outgoing edges are redirected out of the outputs O_v of the ST-robust graph. Because the graph is maximally ST-robust removing k nodes from M_v corresponds to destroying at most $2k$ edges in the original graph G.

Our reduction gives us a fundamentally new way to design node-depth-robust graphs: design an edge-depth-robust graph (easier) and then reduce it to a node-depth-robust graph. The reduction can be used with a construction from [Sch83] to construct a $(\frac{n\log\log n}{\log n}, \frac{n}{\log n (\log n)^{\log\log n}})$-depth-robust graph. We conjecture that several prior DAG constructions (e.g, [EGS75, Sch83, ABP18]) are actually $(n\log\log n, \frac{n}{\log n})$-edge-depth-robust. If any of these conjectures are true then our reduction would immediately yield the desired $(\frac{n\log\log n}{\log n}, \frac{n}{\log n})$-depth-robust graph.

We also present several other applications for maximally ST-robust graphs including the construction of (tight) proofs-of-space and wide block-labeling functions.

2 Edge to Node Depth-Robustness

In this section, we assume the existence of linear sized, constant indegree, maximally ST-robust graphs and use this assumption to construct a transformation of an (e,d)-edge-depth robust graph with m edges into an (e,d)-node-depth robust graph with constant indegree and $O(m)$ nodes. In the next section we will construct a family of ST-robust graphs that satisfies assumption 2.1.

Assumption 2.1. There is a family of graphs $\mathbb{M}=\{M_n\}_{n=1}^{\infty}$ with the property that for each $n\geq 1$, M_n has constant indegree, $O(n)$ nodes, and is maximally ST-robust.

2.1 Reduction Definition

Let $G=(V,E)$ be a DAG, and let $\mathbb{M}$ be as in Assumption 2.1. Then we define Reduce(G, $\mathbb{M}$) in construction 2.2 as follows:

Construction 2.2 (Reduce(G, $\mathbb{M}$)). Let $G=(V,E)$ and let $\mathbb{M}$ be the family of graphs defined above. For each $M_n\in\mathbb{M}$, we say that $M_n=(V(M_n),E(M_n))$, with $V(M_n)=I(M_n)\cup O(M_n)\cup D(M_n)$, where $I(M_n)$ are the inputs of M_n, $O(M_n)$ are the outputs, and $D(M_n)$ are the internal vertices. For $v\in V$, let $\delta(v)=\max\{\text{indegee}(v),\text{outdegree}(v)\}$ Then we define Reduce(G) $=(V_R,E_R)$, where $V_R=\{(v,w)|v\in V, w\in V_{\delta(v)}\}$ and $E_R=E_{internal}\cup E_{external}$. We let $E_{internal}=\{((v,u_{\delta(v)}),(v,w_{\delta(v)}))|v\in V,(u_{\delta(v)},w_{\delta(v)})\in E(H_{\delta(v)})\}$. Then for each $v\in V$, we define an $In(v)=\{u:(u,v)\in E\}$ and $Out(v)=\{u:(v,u)\in E\}$ and then pick two injective mappings $\pi_{in,v}:In(v)\rightarrow I(V_{\delta(v)})$ and $\pi_{out,v}:Out(v)\rightarrow O(V_{\delta(v)})$. We let $E_{external}=\{((u,\pi_{out,u}(v)),(v,\pi_{in,v}(u))) : (u,v)\in E\}$.

Intuitively, to costruct Reduce(G, $\mathbb{M}$) we replace every node of G with a constant indegree, maximally ST-robust graph, mapping the edges connecting two nodes from the outputs of one ST-robust graph to the inputs of another. Then for every $e=(u,w)\in E$, add an edge from an output of $M_{\delta(u)}$ to an input of $M_{\delta(w)}$ such that the outputs of $M_{\delta(u)}$ have outdegree at most 1, and the inputs of $M_{\delta(w)}$ have indegree at most 1. If $v\in V$ is replaced by $M_{\delta(v)}$, then we call v the *genesis node* and $M_{\delta(v)}$ its *metanode*.

2.2 Proof of Main Theorem

We now state the main result of this section which says that if G is edge-depth robust then Reduce(G,$\mathbb{M}$) is node depth-robust.

Theorem 2.3. *Let G be an (e, d)-edge-depth-robust DAG with m edges. Let $\mathbb{M}$ be a family of max ST-Robust graphs with constant indegree. Then $G'=(V',E')=$ Reduce(G,$\mathbb{M}$) is ($e/2$,d)-depth robust. Furthermore, G' has maximum indegree $\max_{v\in V(G)}\{\text{indeg}\left(M_{\delta(v)}\right)\}$, and its number of nodes is $\sum_{v\in V(G)}\left|M_{\delta(v)}\right|$ where $\delta(v)=\max\{\text{indeg}(v),\text{outdeg}(v)\}$.*

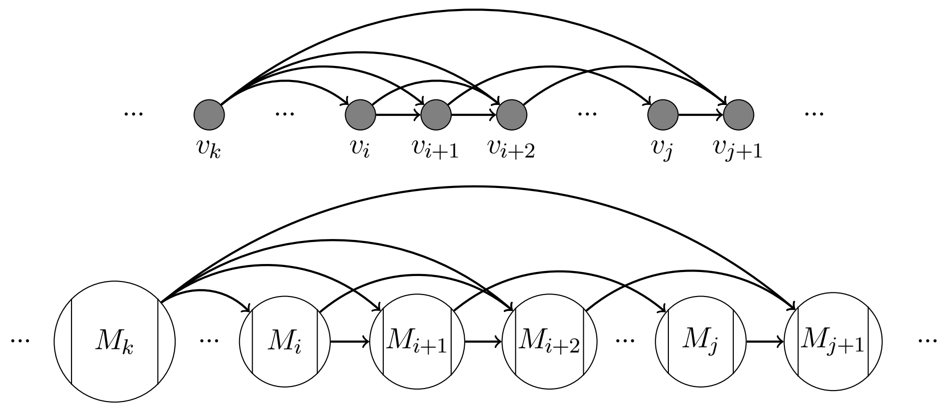

Figure 1: Diagram of the transformation Reduce(G,$\mathbb{M}$)

A formal proof can be found in Appendix B. We briefly outline the intuition for this proof below.

Proof. (Intuition) The first thing we node is that each graph $M_{\delta(v)}$ has constant indegree at most $c\delta(v)$ nodes for some constant $c>0$. Therefore, the graph G' has $\sum_{v\in V(G)}\left|M_{\delta(v)}\right|\leq c\sum_v \delta(v)\leq 2cm$ nodes and G' has constant indegree.

Now for any set $S\subseteq V'$ of nodes we remove from G' we will map S to a corresponding set $S_{irr}\subseteq E$ of at most $|S_{irr}|\leq 2|S|$ irrepairable edges in G. We then prove that any path P in $G-S_{irr}$ corresponds to a longer path P' in $G'-S$ that is at least as long. Intuitively, each incoming edge (u,v) (resp. outgoing edge (v,w)) in $E(G)$ corresponds to an input node (resp. output node) in v's corresponding metanode $M_{\delta(v)}$ which we will label $x_{u,v}$ (resp. $y_{v,w}$). If $S\subseteq V'$ removes at most k nodes from the metanode $M_{\delta(v)}$ then, by maximal ST-robustness, we still can find $\delta(v)-k$ inputs and $\delta(v)-k$ outputs that are *all* pairwise connected. If $x_{u,v}$ (resp. $y_{v,w}$) is not part of this pairwise connected subgraph then we will add the corresponding edge (u,v) (resp. (v,w)) to the set S_{irr}. Thus, the set S_{irr} will have size at most $2|S|$ Claim B.2 in the appendix).

Intuitively, any path P in $G-S_{irr}$ can be mapped to a longer path P' in $G'-S$ (Claim B.1). If P contains the edges (u,v),(v,w) then we know that the input node $x_{u,v}$ and output node $y_{u,v}$ node in $M_{\delta(v)}$ are still connected in $G'-S$. $\square$

Corollary 2.4. *(of Theorem 2.3) If there exists some constants c_1,c_2, such that we have a family $\mathbb{M}=\{M_n\}_{n=1}^{\infty}$ of linear sized $|V(M_n)|\leq c_1 n$, constant indegree* $\text{indeg}(M_n)\leq c_2$, *and maximally ST-robust graphs, then* Reduce(G,$\mathbb{M}$) *has maximum indegree c_2 and the number of nodes is at most $2c_1 m$.*

The next corollary states that if we have a family of maximally ST-robust graphs with $\mathbb{M}=\{M_k\}_{k=1}^{\infty}$ and depth d_k then we can transform any (e,d)-edge-depth-robust DAG $G=(V,E)$ with maximum degree $\delta=\max_{v\in V}\delta(v)$ into $(e/2,d\cdot d_\delta)$-depth robust graph. Instead of replacing each node $v\in G$ with a copy of $M_{\delta(v)}$, we instead replace each node with a copy of $M_{\delta,v}:=M_\delta$, attaching the edges same way as in Construction 2.2. Thus the transformed graph G' has $|V(G)|\times|M_\delta|$ nodes and constant indegree. Intuitively, any path P of length d in $G-S_{irr}$ now maps to a path P' of length $d\times d_\delta$ — if P contains the edges (u,v),(v,w) then we know that the input node $x_{u,v}$ and output node $y_{u,v}$ node in $M_{\delta,v}$ are connected in $G'-S$ *by a path of length at least* d_δ.

Corollary 2.5. *(of Theorem 2.3) Suppose that there exists a family* $\mathbb{M}=\{M_k\}_{k=1}^{\infty}$ *of max ST-robust graphs with depth* d_k *and constant indegree. Given any* (e,d)*-edge-depth-robust DAG* G *with* n *nodes and maximum degree* δ *we can construct a DAG* G' *with* $n\times|M_\delta|$ *nodes and constant indegree that is* $(e/2,d\cdot d_\delta)$*-depth robust.*

Proof. (sketch) Instead of replacing each node $v\in G$ with a copy of $M_{\delta(v)}$, we instead replace each node with a copy of $M_{\delta,v}:=M_\delta$, attaching the edges same way as in Construction 2.2. Thus the transformed graph G' has $|V(G)|\times|M_\delta|$ nodes and constant indegree. Let $S\subset V(G')$ be a set of nodes that we will remove from G'. By Claim B.1, there exists a path P in $G'-S$ that passes through d metanodes $M_{\delta,v_1},\ldots,M_{\delta,v_d}$. Since M_δ is maximally ST-robust with depth d_δ the sub-path $P_i=P\cap M_{\delta,v_i}$ through each metanode has length $|P_i|\geq d_\delta$. Thus, the total length of the path is at least $\sum_i|P_i|\geq d\cdot d_\delta$. □

Corollary 2.6. *(of Theorem 2.3) Let* $\epsilon>0$ *be any fixed constant. Given any family* $\{G_m\}_{m=1}^{\infty}$ *of* (e_m,d_m)*-edge-depth-robust DAGs* G_m *with* m *nodes and maximum indegree* δ_m *then for some constants* $c_1,c_2>0$ *we can construct a family* $\{H_m\}_{m=1}^{\infty}$ *of DAGs such that each DAG* H_m *is* $(e_m/2,d_m\cdot\delta_m^{1-\epsilon})$*-depth robust,* H_m *has maximum indegree at most* c_2 *(constant) and at most* $\left|V(H_m)\right|\leq c_1 m\delta_m$ *nodes.*

Proof. (sketch) This follows immediately from Corollary 2.5 and from our construction of a family $\mathbb{M}_\epsilon=\{M_{k,\epsilon}\}_{k=1}^{\infty}$ of max ST-robust graphs with depth $d_k>k^{1-\epsilon}$ and constant indegree. □

Corollary 2.7. *(of Theorem 2.3) Let* $\{e_m\}_{m=1}^{\infty}$ *and* $\{d_m\}_{m=1}^{\infty}$ *be any sequence. If there exists a family* $\{G_m\}_{m=1}^{\infty}$ *of* (e_m,d_m)*-edge-depth-robust graphs, where each DAG* G_m *has* m *edges, then there is a corresponding family* $\{H_n\}_{n=1}^{\infty}$ *of constant indegree DAGs such that each* H_n *has* n *nodes and is* $(\Omega(e_n),\Omega(d_n))$*-depth-robust.*

The original Grate's construction [Sch83], G, has $N=2^n$ nodes and $m=n2^n$ edges and for any $s\leq n$, and is $(s2^n, \frac{N}{\sum_{j=0}^{s}\binom{n}{j}})$-edge-depth-robust. For node depth-robustness we only had matching constructions when $s=O(1)$ [ABP17, ABH17] and $s=\Omega(\log N)$ [Sch83] — no comparable lower bounds were known for intermediate s.

Corollary 2.8. *(of Theorem 2.3) There is a family of constant indegree graphs $\{G_n\}$ such that G_n has $O(N=2^n)$ nodes and G_n is $(sN/(2n), \frac{N}{\sum_{j=0}^{s}\binom{n}{j}})$-edge-depth-robust for any $1\leq s\leq \log n$*

In particular, setting $s=\log\log n$ and applying the indegree reduction from Theorem 2.3, we see that the transformed graph G' has constant indegree, $N'=O(n2^n)$ nodes, and is $(\frac{N'\log\log N'}{\log N'}, \frac{N'}{\log N'(\log N')^{\log\log N'}})$-depth-robust. Blocki et al. [BHK+19] showed that if there exists a node depth robust graph with $e=\Omega(N\log\log N/\log N)$ and $d=\Omega(N\log\log N/\log N)$ then one can obtain another constant indegree graph with pebbling cost $\Omega(N^2\log\log N/\log N)$ which is optimal for constant indegree graphs. We conjecture that the graphs in [EGS75] are sufficiently edge depth robust to meet these bounds after being transformed by our reduction.

3 ST Robustness

In this section we show how to construct maximally ST-robust graphs with constant indegree and linear size. We first introduce some of the technical building blocks used in our construction including superconcentrators [Val76, Pip77, GG81] and grates [Sch83]. Using these building blocks we then provide a randomized construction of a c_1-maximally ST-robust DAG with linear size and constant indegree for some constant $c_1>0$ — sampled graphs are c_1-maximally ST-robust DAG with high probability. Finally, we use c_1-maximally ST-robust DAGs to construct a family of maximally ST-robust graphs with linear size and constant indegree.

3.1 Technical Ingredients

We now introduce other graph properties that will be useful for constructing ST-robust graphs.

Grates A DAG $G=(V,E)$ with n inputs I and n outputs O is called a (c_0,c_1)-grate if for any subset $S\subset V$ of size $|S|\leq c_0 n$ at least $c_1 n^2$ input output pairs $(x,y)\in I\times O$ remain connected by a directed path from x to y in $G-S$. Schnitger [Sch83] showed how to construct (c_0,c_1)-grates with $O(n)$ nodes and constant indegree for suitable constants $c_0,c_1>0$. The notion of an maximally ST-robust graph is a strictly stronger requirement since there is no requirement on which pairs are connected. However, we show that a slight modification of Schnitger's [Sch83] construction is a $(cn,n/2)$-ST-robust for a suitable constant c. We then transform this graph into a c_1-maximally ST-robust graph by sandwiching it in between two superconcentrators. Finally, we show how to use several c_1-maximally ST-robust graphs to construct a maximally ST-robust graph.

Superconcentrators We say that a directed acyclic graph $G=(V,E)$ with n input vertices and n output vertices is an ***n*-superconcentrator** if for any r inputs and any r outputs, $1\le r\le n$, there are r vertex-disjoint paths in G connecting the set of these r inputs to these r outputs. We note that there exists linear size, constant indegree superconcentrators [Val76, Pip77, GG81] and we use this fact throughout the rest of the paper. For example, Pippenger [Pip77] constructed an n-superconcentrator with at most $41n$ vertices and indegree at most 16.

Connectors We say that an n-superconcentrator is an ***n*-connector** if it is possible to specify which input is to be connected to which output by vertex disjoint paths in the subsets of r inputs and r outputs. Connectors and superconcentrators are potential candidates for ST-robust graphs because of their highly connective properties. In fact, we can prove that any connectors ***n*-connector** is maximally ST-robust — the proof of Theorem 3.1 can be found in the appendix. While we have constructions of ***n*-connector** graphs these graphs have $O(n\log n)$ vertices and indegree of 2, an information theoretic technique of Shannon [Sha50] can be used to prove that any n-connector with constant indegree requires *at least* $\Omega(n\log n)$ vertices — see discussion in the appendix. Thus, we cannot use n-connectors to build linear sized ST-robust graphs. However, Shannon's information theoretic argument does not rule out the existence of linear size ST-robust graphs.

Theorem 3.1. *If G is an n-connector, then G is $(k,n-k)$-ST-robust, for all $1\le k\le n$.*

3.2 Linear Size ST-robust Graphs

ST-robust graphs have similar connective properties to connectors, so a natural question to ask is whether ST-robust graphs with constant indegree require $\Omega(n\log n)$ vertices. In this section, we show that linear size ST-robust graphs exist by showing that a modified version of the Grates construction [Sch83] becomes c-maximally ST-robust when sandwiched between two superconcentrators for some constant c.

In the proof of Theorem A in [Sch83], Schnitger constructs a family of DAGs $(H_n | n \in N)$ with constant indegree δ_H, where n is the number of nodes and H_n is $(cn, n^{2/3})$-depth-robust, for suitable constant $c > 0$. We construct a similar graph G_n as follows:

Construction 3.2 (G_n)**.** We begin with H_n^1, H_n^2 and H_n^3, three isomorphic copies of H_n with disjoint vertex sets V_1, V_2 and V_3. For each top vertex $v \in V_3$ sample τ vertices $x_1^v, \ldots, x_\tau^v$ independently and uniformly at random from V_2 and for each $i \leq \tau$ add each directed edge (x_i^v, v) to G_n to connect each of these sampled nodes to v. Similarly, for each node vertex $u \in V_2$ sample τ vertices $x_1^u, \ldots, x_\tau^u$ from V_1 independently and uniformly at random and add each directed edge (x_i^u, u) to G_n. Note that $\mathsf{indeg}(G_n) \leq \mathsf{indeg}(H_n) + \tau$.

Schnitger's construction only utilizes two isomorphic copies of H_n and the edges connecting H_n^1 and H_n^2 a sampled by picking τ random permutations. In our case the analysis is greatly simplified by picking the edges uniformly and we will need three layers to prove ST-robustness. We will use the following lemma from the Grates paper as a building block. A proof of Lemma 3.3 is included in the appendix for completeness.

Lemma 3.3. *[Sch83] For some suitable constant $c > 0$ any any subset S of $cn/2$ vertices of G_n the graph $H_n^1 - S$ contains $k = cn^{1/3}/2$ vertex disjoint paths $A_1, \ldots, A_k$ of length $n^{2/3}$ and $H_n^2 - S$ contains k vertex disjoint paths $B_1, \ldots, B_k$ of the same length.*

We use Lemma 3.3 to help establish our main technical Lemma 3.4. We sketch the proof of Lemma 3.4 below. A formal proof can be found in Appendix B.

Lemma 3.4. *Let G_n be defined as in Construction 3.2. Then for some constants $c > 0$, with high probability G_n has the property that for all $S \subset V(G_n)$ with $|S| = cn/2$ there exists $A \subseteq V(H_n^1)$ and $B \subseteq V(H_n^3)$ such that for every pair of nodes $u \in A$ and $v \in B$ the graph $G_n - S$ contains a path from u to v and $|A|, |B| \geq 9cn/40$.*

Proof. (Sketch) Fixing any S we can apply Lemma 3.3 to find $k := cn^{1/3}/2$ vertex disjoint paths $P_{1,S}^i, \ldots, P_{k,S}^i$ in H_n^i of length $n^{2/3}$ for each $i \leq 3$. Here, c is the constant from Lemma 3.3. Let $U_{j,S}^i$ be the upper half of the j-th path in H_n^i and $L_{j,S}^i$ be the lower half and define the event $BAD_{i,S}^{upper}$ to be the event that there exists at least $k/10$ indices $j \leq k$ s.t., $U_{j,S}^2$ is disconnected from $L_{i,S}^3$. We construct B by taking the union of all of upper paths $U_{i,S}^3$ in H_n^3 for each non-bad (upper) indices i. Similarly, we define $BAD_{i,S}^{lower}$ to be the event that there exists at least $k/10$ indices $j \leq k$ s.t. $U_{i,S}^1$ is disconnected from $L_{j,S}^2$ and we construct A

be taking the union of all of the lower paths $L^1_{i,S}$ in H^1_n for each non-bad (lower) indices i. We can now argue that any pair of nodes $u\in A$ and $v\in B$ is connected by invoking the pigeonhole principle i.e., if $u\in L^1_{i,S}$ and $v\in U^3_{i',S}$ for good indices i and i' then there exists some path P^2_j in the middle layer H^2_n which can be used to connect u to v. We still need to argue that $|A|,|B|\geq cn/3$ for some constant c. To lower bound $|B|$ we introduce the event $BAD_S=|\{i : BAD^{upper}_{i,S}\}|>\frac{k}{10}$ and note that unless BAD_S occurs we have $|B|\geq(9k/10)n^{2/3}/2=9cn/40$. Finally, we show that $\mathbb{P}[BAD_S]$ is very small and then use union bounds to show that, for a suitable constant τ, the probability $\mathbb{P}[\exists S BAD_S]$ becomes negligibly small. A symmetric argument can be used to show that $|A|\geq 9cn/40$. □

We now use G_n to construct c-maximally ST-robust graphs with linear size.

Construction 3.5 (M_n)**.** We construct the family of graphs M_n as follows: Let the graphs SC^1_n and SC^2_n be linear sized n-superconcentrators with constant indegree δ_{SC} [Pip77], and let H^1_n, H^2_n and H^3_n be defined and connected as in G_n, where every output of SC^1_n is connected to a node in H^1_n, every node of H^3_n is connected to an input of SC^2_n.

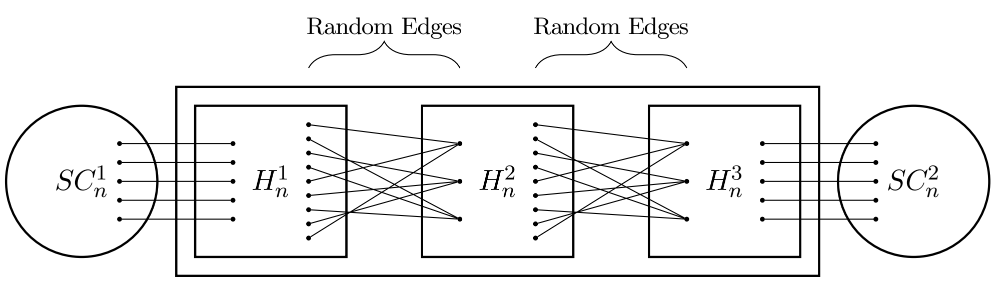

Figure 2: A diagram of the constant indegree, linear sized, ST-robust graph M_n.

Theorem 3.6. *There exists a constant $c'>0$ such that for all sets $S\subset V(M_n)$ with $|S|\leq c'n/2$, M_n is $(|S|,n-|S|)$-ST-robust, with n inputs and n outputs and constant indegree.*

Proof. Let $c'=9c/40$, where c is the constant from G_n. Consider M_n-S. Then because $|S\cap(H^1_n\cup H^2_n)|\leq|S|\leq c'n/2\leq cn/2$, by Lemma 3.4 with a high probability there exists sets A in H^1_n and B in H^3_n with $|A|,|B|\geq\frac{9}{10}k\frac{n^{2/3}}{2}=\frac{9}{40}cn=c'n$, such that every node in A connects to every node in B. By the properties of superconcentrators, the size of the set BAD_1 of inputs u in SC^1_n that can't reach *any* node in A in M_n-S. We claim that $|BAD_1|\leq|S|\leq c'n$. Assume for contradiction that

$|BAD_1|>|S|$ then SC_n^1 contains at least $\min\{|BAD_1|,|A|\}>|S|$ node disjoint paths between BAD_1 and A. At least one of these node disjoint paths does not intersect S which contradicts the definition of BAD_1. Similarly, we can bound the size of BAD_2, the set of outputs in SC^n which are not reachable from any node in B. Given any input $u \notin BAD_1$ of SC_n^1 and any output $v \notin BAD_2$ of SC_n^2 we can argue that u is connected to v in $M_n - S$ since we can reach some node $x \in A$ from u and v can be reached from some node $y \in B$ and *any* such pair x,y is connected by a path in $M_n - S$. It follows that M_n is $(|S|, n-|S|)$-ST-robust. □

Corollary 3.7. *(of Theorem 3.6) For all $\epsilon > 0$, there exists a family of DAGs $\mathbb{M} = \{M_n^\epsilon\}_{n=1}^\infty$, where each M_n^ϵ is a c-maximally ST-robust graphs with $|V(M_n)| \le c_\epsilon n$, indegree$(M_n) \le c_\epsilon$, and depth $d = n^{1-\epsilon}$.*

Proof. (sketch) In the proof of Lemma 3.3, we used $(cn, n^{2/3})$-depth robust graphs. When considering the paths A_i and B_j, we were considering connecting the upper half of one path to the lower half of another. Thus, after we remove nodes from M_n, there exists a path of length at least $n^{2/3}$ that connects any remaining input to any remaining output. Thus M_n is c-maximally ST-robust with depth $d = n^{2/3}$. In [Sch83], Schnitger provides a construction that is $(cn, n^{1-\epsilon})$-depth robust for all constant $\epsilon > 0$. By the same arguments we used in this section, we can construct c-maximally ST-robust graphs with depth $d = n^{1-\epsilon}$, where the constant c depends on ϵ. □

3.3 Constructing Maximal ST-Robust Graphs

In this section, we construct maximal ST-robust graphs, which are 1-maximally ST-robust, from c-maximally ST-robust graphs. We give the following construction:

Construction 3.8 ($\mathbb{O}(M_n)$)**.** Let M_n be a c-maximally ST-robust graph on $O(n)$ nodes. Let O be a set $o_1, o_2, ..., o_n$ of n output nodes and let I be a set $i_1, i_2, ..., i_n$ of n input nodes. Let S_j for $1 \le j \le \lceil \frac{1}{c} \rceil$ be a copy of M_n with outputs $o_1^j, o_2^j, ..., o_n^j$ and inputs $i_1^j, i_2^j, ..., i_n^j$. Then for all $1 \le j \le n$ and for all $1 \le k \le n$, add a directed edge from i_k to i_k^j and from o_k^j to o_k.

Because we connect $\lceil \frac{1}{c} \rceil$ copies of M_n to the output nodes, $\mathbb{O}(M_n)$ has indegree $\max\{\delta, \lceil \frac{1}{c} \rceil\}$, where δ is the indegree of M_n. Also, if M_n has kn nodes, then $\mathbb{O}(M_n)$ has $(k\lceil \frac{1}{c} \rceil + 2)n$ nodes. We now show that $\mathbb{O}(M_n)$ is a maximal ST-robust graph.

Theorem 3.9. *Let M_n be a c-maximally ST-robust graph. Then $\mathbb{O}(M_n)$ is a maximal ST-robust graph.*

Proof. Let $R \subset V(\mathbb{O}(M_n))$ with $|R| = k$. Let $R = R_I \cup R_M \cup R_O$, where $R_I = R \cap I$, $R_O = R \cap O$, and $R_M = R \cap \left(\cup_{i=1}^{\lceil 1/c \rceil} S_i\right)$. Consider $\mathbb{O}(M_n) - R$. We see that $|R_M| \le k$, so by the Pidgeonhole Principal at least one S_j has less than cn nodes removed, say it has t nodes removed for $t \le cn$. Hence $t \le |R_M|$. Since S_j is c-max ST-robust there exists a subgraph H of S_j R containing $n-t$ inputs and $n-t$ outputs such that every input is connected to all of the outputs. Let H' be the subgraph induced by the nodes in $V(H) \cup I' \cup O'$, where $I' = \{(i_a, i_a^b) | i_a^b \in H\}$ and $O' = \{(o_a^b, o_a) | o_a^b \in H\}$.

Claim 3.10. The graph H' contains at least $n-k$ inputs and $n-k$ outputs and there is a path between every pair of input and output nodes.

Proof. The set $|I \setminus I'| \le |I \cap R| + |V(S_j) \cap R| \le |R| \le k$. Similarly, $|O \setminus O'| \le |O \cap R| + |V(S_j) \cap R| \le |R| \le k$. Let $v \in I'$ be some input. By the connectivity of H, v can reach all of the outputs in O'. Thus there is a path between every pair of input and output nodes. □

Thus $\mathbb{O}(M_n)$ is $(k, n-k)$-ST-robust for all $1 \le k \le n$. Therefore $\mathbb{O}(M_n)$ is a maximal ST-robust graph. □

Corollary 3.11. *(of Theorem 3.9) For all $\epsilon > 0$, there exists a family $\mathbb{M}^\epsilon = \{M_k^\epsilon\}_{k=1}^\infty$ of max ST-robust graphs of depth $d = n^{1-\epsilon}$ such that $|V(M_k^\epsilon)| \le c_\epsilon n$ and indegree$(M_k^\epsilon) \le c_\epsilon$.*

Proof. Apply Construction 3.8 to the family graphs $\mathbb{M}^\epsilon = \{M_k^\epsilon\}_{k=1}^\infty$ from Corollary 3.7. Then by Theorem 3.9, the family of graphs $\{\mathbb{O}(M_k^\epsilon)\}_{k=1}^\infty$ is the desired family. □

4 Applications of ST-Robust Graphs

As outlined previously maximally ST-Robust graphs give us a tight connection between edge-depth robustness and node-depth robustness. Because edge-depth-robust graphs are often easier to design than node-depth robust graphs [Sch83] this gives us a fundamentally new approach to construct node-depth-robust graphs. Beyond this exciting connection we can also use ST-robust graphs to construct perfectly tight proofs of space [Pie19,Fis19] and asymptotically superior wide-block labeling functions [CT19].

4.1 Tight Proofs of Space

In Proof of Space constructions [Pie19] we want to find a DAG $G=(V,E)$ with small indegree along with a challenge set $V_C \subseteq V$. Intuitively, the prover will label the graph G using a hash function H (often modeled as a random oracle in security proofs) such that a node v with parents $v_1,...,v_\delta$ is assigned the label $L_v = H(L_{v_1},...,L_{v_\delta})$. The prover commits to storing L_v for each node v in the challenge set V_C. The pair (G,V_C) is said to be (s,t,ϵ)-hard if for any subset $S \subseteq V$ of size $|S| \leq s$ at least $(1-\epsilon)$ fraction of the nodes in V_C have depth $\geq t$ in $G-S$ — a node v has depth $\geq t$ in $G-S$ if there is a path of length $\geq t$ ending at node v. Intuitively, this means that if a cheating prover only stores $s \leq |V_C|$ labels and is challenged to reproduce a random label L_v with $v \in V_C$ that, except with probability ϵ, the prover will need at least t sequential computations to recover L_v — as long as t is sufficiently large the verifier the cheating prover will be caught as he will not be able to recover the label L_v in a timely fashion. Pietrzak argued that (s,t,ϵ)-hard graphs translate to secure Proofs of Space in the parallel random oracle model [Pie19].

We want G to have small indegree $\delta(G)$ (preferably constant) as the prover will need $O(N\delta(G))$ steps. Additionally, we want $|V_C| = \Omega(N)$ and ϵ to be small while s,t should be larger. Pietrzak [Pie19] proposed to let G_ϵ be an ϵ-extreme depth-robust graph with $N' = 4N$ nodes and to let $V_C = [3N+1, 4N]$ be the last N nodes in this graph. An ϵ-extreme depth-robust graph with N' nodes is (e,d)-depth robust for any $e+d \leq (1-\epsilon)N$. Such a graph is $(s,N,s/N+4\epsilon)$-hard for any $s \leq N$. Alwen et al. [ABP18] constructed ϵ-extreme depth-robust graphs with indegree $\delta(G) = O(\log N)$ though the hidden constants seem to be quite large. Thus, it would take time $O(N \log N)$ for the prover to label the graph G. We remark that $\epsilon = s/|V_C|$ is the tightest possible bound one can hope for as the prover can always store s labels from the set V_C.

We remark that if we take V_C to be any subset of output nodes from a maximally ST-robust graph and overlay an $(e=s, d=t)$-depth robust graph over the input nodes, then the resulting graph will be $(s,t,\epsilon = s/|V_C|)$-hard — optimally tight in ϵ. In particular, given a DAG $G=(V=[N],E)$ with N nodes devine the overlay graph H_G by starting with a maximally ST-robust graph with $|V|$ inputs $I = \{x_1,...,x_{|V|}\}$ and $|V|$ outputs O then for every directed edge $(u,v) \in E(G)$ we add the directed edge (x_u,x_v) to $E(H_G)$ and we specify a target set $V_C \subseteq O$. Fisch [Fis19] gave a practical construction of (G,V_C) with indegree $O(\log N)$ that is $(s,N,\epsilon = s/N+\epsilon')$-hard. The constant ϵ' can be arbitrarily small though the number of nodes in the graph scales with $O(N \log 1/\epsilon')$. Utililizing ST-robust graphs we fix $\epsilon' = 0$ without increasing the size of the graph[1].

[1] As a disclaimer we are not claiming that our construction would be more efficient than

Theorem 4.1. *If G is (e,d)-depth robust then the pair (H_G,V_C) specified above is $(s,t=d+1,s/|V_C|)$-hard for any $s \leq e$.*

Proof. Let S be a subset of $|S| \leq s$ nodes in H_G. By maximal ST-robustness we can find a set A of $N-|S|$ inputs and B of $N-|S|$ outputs such that every pair of nodes $u \in A$ and $v \in B$ are connected in $H_G - S$. We also note since A contains all but s nodes from G that some node $u \in A$ is the endpoint of a path of length t by (s,t)-depth-robustness of G. Since u is connected to *every* node in B this means that every node $v \in B$ is the endpoint of a path of length *at least* $t+1$. □

This result immediately leads to a $(s,N^{1-\epsilon},s/N)$-hard pair for any $s \leq N$ which the prover can label in $O(N)$ time as the DAG G has constant indegreee. We expect that in many settings $t=N^{1-\epsilon}$ would be sufficiently large to ensure that a cheating prover is caught with probability s/N after each challenge i.e., if the verifier expects a response within 3 seconds, but it would take longer to evaluate the hash function H a total of $N^{1-\epsilon}$ sequential times.

Corollary 4.2. *For any constant $\epsilon > 0$ there is a constant indegree DAG G with $O(N)$ nodes along with a target set $V_C \subseteq V(G)$ of N nodes such that the pair (G,V_C) is $(s,t=N^{1-\epsilon},s/N)$-hard for any $s \leq N$.*

Proof. (sketch) Let G be an $\left(N,N^{1-\epsilon}\right)$-depth robust graph with $N' = O(N)$ nodes and constant indegree from [Sch83]. We can then take V_C to be any subset of N output nodes in the graph H_G and apply Theorem 4.1. □

If one does not want to relax the requirement that $t=\Omega(N)$ then we can provide a perfectly tight construction with $O(N\log N)$ nodes and constant indegree. Since the graph has constant indegree it will take $O(N\log N)$ work for the prover to label the graph. This is equivalent to [Pie19], but with perfect tightness $\epsilon = s/N$.

Corollary 4.3. *For any constant $\epsilon > 0$ there is a constant indegree DAG G with $N' = O(N\log N)$ nodes along with a target set $V_C \subseteq V(G)$ of N nodes such that the pair (G,V_C) is $(s,t,s/N)$-hard for any $s \leq N$.*

Proof. (sketch) Let G be an $(N,N\log N)$-depth robust graph with $N' = O(N\log N)$ nodes and constant indegree from [ABH17]. We can then take V_C to be any subset of N output nodes in the graph H_G and apply Theorem 4.1. □

Finally, if we want to ensure that the graph only has $O(N)$ nodes and $t=\Omega(N)$ we can obtain a perfectly tight construction with indegree $\delta(G) = O(\log N)$.

[Fis19] for practical parameter settings.

Corollary 4.4. *For any constant $\epsilon > 0$ there is a DAG G with $O(N)$ nodes and indegree $\delta(G) = O(\log N)$ along with a target set $V_C \subseteq V(G)$ of N nodes such that the pair (G, V_C) is $(s, N, s/N)$-hard for any $s \leq N$.*

Proof. (sketch) Let G be an (N,N)-depth robust graph with $N' = 3N$ nodes from [ABP18]. We can then take V_C to be any subset of N output nodes in the graph H_G and apply Theorem 4.1. □

4.2 Wide-Block Labeling Functions

Chen and Tessaro [CT19] introduced source-to-sink depth robust graphs as a generic way of obtaining a wide-block labeling function $H_{\delta,W} : \{0,1\}^{\delta W} \to \{0,1\}^W$ from a small-block function $H_{fix} : \{0,1\}^{2L} \to \{0,1\}^L$ (modeled as an ideal primitive). In their proposed approach one transforms a graph G with indegree δ and into a new graph G' by replacing every node with a source-to-sink depth-robust graph. Labeling a graph G with a wide-block labeling function is now equivalent to labeling G' with the original labeling function H_{fix}. The formal definition of Source-to-Sink-Depth-Robustness is presented below:

Definition 4.5 (Source-to-Sink-Depth-Robustness (SSDR) [CT19]). A DAG $G = (V, E)$ is (e,d)-source-to-sink-depth-robust (SSDR) if and only if for any $S \subset V$ where $|S| \leq e$, $G - S$ has a path (with length at least d) that starts from a source node of G and ends up in a sink node of G.

If G is (e,d)-depth robust and G' is constructed by replacing every node v in G with a (e^*, d^*)-source-to-sink-depth-robust (SSDR) and orienting incoming (resp. outgoing) edges into the sources (resp. out of the sinks) then the graph G' is (ee^*, dd^*)-depth robust [CT19] and has cumulative pebbling complexity at least $ed(e^* d^*)$ [ABP17]. The SSDR graphs constructed in [CT19] are $(\frac{K}{4}, \frac{\delta K^2}{2})$-SSDR with $O(\delta K^2)$ vertices and constant indegree. They fix $K := W/L$ as the ratio between the length of outputs for $H_{\delta,W} : \{0,1\}^{\delta W} \to \{0,1\}^W$ and the ideal primitive H_{fix}. Their graph has δK source nodes for a tunable parameter $\delta \in \mathbb{N}$, $O(\delta K^2)$ vertices and constant indegree. Ideally, since we are increasing the number of nodes by a factor of δK^2 we would like to see the cumulative pebbling complexity increase by a quadratic factor of $\delta^2 K^4$. Instead, if we start with an (e,d)-depth robust graph with cumulative pebbling complexity $O(ed)$ their final graph G' has cumulative pebbling complexity $ed \times \frac{\delta K^3}{8}$. Chen and Tessaro left the problem of finding improved source-to-sink depth-robust graphs as an open research question.

Our construction of ST-robust graphs can asymptotically[2] improve some of

[2] While we improve the asymptotic performance we do not claim to be more efficient for practical values of δ, K.

their constructions, specifically their constructions of source-to-sink-depth-robust graphs and wide-block labeling functions.

Theorem 4.6. *Let G be a maximal ST-robust graph with depth d and n inputs and outputs. Then G is an $(n-1,d)$-SSDR graph.*

Proof. By the maximal ST-robustness property, $n-1$ arbitrary nodes can be removed from G and there will still exist at least one input node that is connected to at least one output node. Since G has depth d, the path between the input node and output node must have length at least d. $\square$

By applying Theorem 4.6 to the construction in Corollary 3.9, we can construct a family of $(\delta K,(\delta K)^{1-\epsilon})$-SSDR graphs with $O(\delta K)$ nodes and constant indegree and δK sources. In this case the cumulative pebbling complexity of our construction would be already be $ed \times \delta^2 K^{2-\epsilon}$ which is much closer to the quadratic scaling that we would ideally like to see. We are off by just K^ϵ for a constant $\epsilon > 0$ that can be arbitrarily small. To make the comparison easier we could also applying Theorem 4.6 to obtain a family of $(\delta K^2,(\delta K^2)^{1-\epsilon})$-SSDR graphs with $O(\delta K^2)$-nodes and constant indegree. While the size of the SSDR matches [CT19] our new graph is $(e\delta K^2,d(\delta K^2)^{1-\epsilon})$-depth robust and has cumulative pebbling complexity $ed \times \delta^{2-\epsilon} K^{4-2\epsilon} \gg ed\delta K^3$.

References

[AB16] Joël Alwen and Jeremiah Blocki. Efficiently computing data-independent memory-hard functions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 241–271. Springer, Heidelberg, August 2016.

[ABH17] Joël Alwen, Jeremiah Blocki, and Ben Harsha. Practical graphs for optimal side-channel resistant memory-hard functions. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1001–1017. ACM Press, October / November 2017.

[ABP17] Joël Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. Depth-robust graphs and their cumulative memory complexity. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 3–32. Springer, Heidelberg, April / May 2017.

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

[BHK+19] Jeremiah Blocki, Benjamin Harsha, Siteng Kang, Seunghoon Lee, Lu Xing, and Samson Zhou. Data-independent memory hard functions: New attacks and stronger constructions. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 573–607. Springer, Heidelberg, August 2019.

[CT19] Binyi Chen and Stefano Tessaro. Memory-hard functions from cryptographic primitives. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 543–572. Springer, Heidelberg, August 2019.

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

[EGS75] Paul Erdös, Ronald L. Graham, and Endre Szemeredi. On sparse graphs with dense long paths., 1975.

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

[GG81] Ofer Gabber and Zvi Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981.

[MMV13] Mohammad Mahmoody, Tal Moran, and Salil P. Vadhan. Publicly verifiable proofs of sequential work. In Robert D. Kleinberg, editor, *ITCS 2013*, pages 373–388. ACM, January 2013.

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

[Pip77] Nicholas Pippenger. Superconcentrators. *SIAM J. Comput.*, 6(2):298–304, 1977.

[Sch83] Georg Schnitger. On depth-reduction and grates. In *24th Annual Symposium on Foundations of Computer Science, Tucson, Arizona, USA, 7-9 November 1983*, pages 323–328. IEEE Computer Society, 1983.

[Sha50] Claude Shannon. Memory requirements in a telephone exchange. *Bell System Technical Journal*, 29(3):343–349, July 1950.

[Val76] Leslie G. Valiant. Graph-theoretic properties in computational complexity. *J. Comput. Syst. Sci.*, 13(3):278–285, December 1976.

[Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In Jozef Gruska, editor, *Mathematical Foundations of Computer Science 1977, 6th Symposium, Tatranska Lomnica, Czechoslovakia, September 5-9, 1977, Proceedings*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977.

Appendix

A Connector Graphs

We say that a directed acyclic graph $G=(V,E)$ with n input vertices and n output vertices is an ***n*-connector** if for any *ordered list* $x_1,...,x_r$ of r inputs and any *ordered list* $y_1,...,y_r$ of r outputs, $1 \le r \le n$, there are r vertex-disjoint paths in G connecting input node x_i to output node y_i for each $i \le r$.

A.1 Connector Graphs are ST-Robust

We remarked in the paper that any ***n*-connector** is maximally ST-robust.
Reminder of Theorem 3.1. *If G is an n-connector, then G is $(k,n-k)$-ST-robust, for all $1 \le k \le n$. Proof of Theorem 3.1.* Let $D \subseteq V(G)$ with $|D|=k$. Consider $G-D$. Let $A=\{(s_1,t_1),...,(s_m,t_m)\}$, where the input $s_i \in S$ is disconnected from the output $t_i \in T$ in $G-D$, or $s_i \in D$ or $t_i \in D$. Let $B=\emptyset$.

Perform the following procedure on A and B: Pick any pair $(s_p,t_p) \in A$ and add s_p and t_p to B. Then remove the pair from A along with any other pair in A that shares either s_p or t_p. Continue until A is empty.

If we consider the nodes of B in G, then there are $|B|$ vertex-disjoint paths between the pairs in B by the connector property, and in $G-D$ at least one vertex is removed from each path. Thus $|B| \le k$, or we have a contradiction.

If $(s,t) \in G-(D \cup B)$ are an input to output pair, and s is disconnected from t, then by the definition of A and B we would have a contradiction, since

(s,t) would still be in A. Thus all of the remaining inputs in $G-(D\cup B)$ are connected to all the remaining outputs.

Hence, if we let $H=G-(D\cup B)$, then H is a subgraph of G with at least $n-k$ inputs and $n-k$ outputs, and there is a path going from each input of H to each of its outputs. Therefore, G is $(k,n-k)$-ST-robust for all $1\le k\le n$. $\square$

Butterfly Graphs A well known family of constant indegree n-connectors, for $n=2^k$, are the k-dimensional butterfly graphs B_k, which are formed by connecting two FFT graphs on n inputs back to back. See Figure A.1 for an example. By Theorem 3.1, the butterfly graph is also a maximally ST-robust graph. However, the butterfly graph has $\Omega(n\log n)$ nodes and does not yield a ST-robust graph of linear size. Since B_k has $O(n\log n)$ vertices and indegree of 2, a natural question to ask is if there exists n-connectors with $O(n)$ vertices and constant indegree.

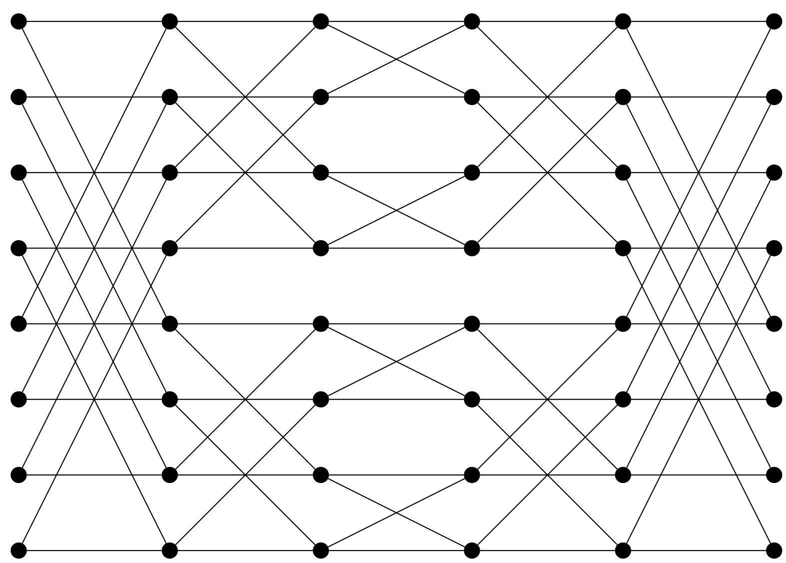

Figure 3: The butterfly graph B_3 is both an 8-superconcentrator and an 8-connector. All edges are directed from left to right.

A.2 Connector Graphs Have $\Omega(n \log n)$ vertices

An information theoretic argument of Shannon [Sha50] rules out the possibility of linear size n-connectors.

Theorem A.1. *(Shannon [Sha50]) An n-connector with constant indegree requires at least $\Omega(n\log n)$ vertices.*

Intuitively, given a n-connector with constant indegree with constant indegree and m edges Shannon argued that we can use the n-connector to encode any permutatation of $[n]$ using m bits. In more detail fixing any permuation π we can find n node disjoint paths from input i to output $\pi(i)$. Because the paths are node disjoint we can encode π simply by specifying the subset S_π of directed edges which appear in one of these node disjoint paths. We require at most m bits to encode S_π and from S_π we can reconstruct the set of node-disjoint paths and recover π. Thus, we must have $m=\Theta(n\log n)$ since we require $\log n!=\Theta(n\log n)$ bits to encode a permutation.

We stress that this information theoretic argument breaks down if the graph G is only ST-robust. We are guaranteed that G contains a path from input i to output $\pi(i)$, but we are not guaranteed that all of the paths are node disjoint. Thus, S_π is insufficient to reconstruct π.

B Missing Proofs

Reminder of Theorem 2.3. *Let G be an (e, d)-edge-depth-robust DAG with m edges. Let $\mathbb{M}$ be a family of max ST-Robust graphs with constant indegree. Then $G'=(V',E')=\mathrm{Reduce}(G,\mathbb{M})$ is $(e/2,d)$-depth robust. Furthermore, G' has maximum indegree $\max_{v\in V(G)}\{\mathrm{indeg}\left(M_{\delta(v)}\right)\}$, and its number of nodes is $\sum_{v\in V(G)}\left|M_{\delta(v)}\right|$ where $\delta(v)=\max\{\mathrm{indeg}(v),\mathrm{outdeg}(v)\}$.*

Proof of Theorem 2.3.

We know that each graph in $\mathbb{M}$ has constant indegree, and that each node v in G will be replaced with a graph in $\mathbb{M}$ with indegree $\mathrm{indeg}\left(M_{\delta(v)}\right)$. Thus G' has maximum indegree $\max_{v\in V(G)}\{\mathrm{indeg}\left(M_{\delta(v)}\right)\}$. Furthermore, the metanode corresponding to the node v has size $|M_{\delta(v)}|$. Thus G' has $\sum_{v\in V(G)}\left|M_{\delta(v)}\right|$ nodes.

Let $S\subset V(G')$ be a set of nodes that we will remove from G'. For a specific node $v\in V(G)$ we let $S_v=S\cap(\{v\}\times V_{\delta(v)})$ denote the subset of nodes deleted from the corresponding metanode. We say that the node $v\in V(G)$ is *irrepairable* with respect to S if $|S_v|\geq\delta(v)$; otherwise we say that v is repairable. If a node v is repairable, then because the metanodes are maximally ST-Robust we can find subsets $I_{v,S}$ and $O_{v,S}$ (with $|I_{v,S}|,|O_{v,S}|\geq\delta(v)-|S_v|$) such that each input node $s\in I_{v,S}$ is connected to every output node in $O_{v,S}$.

We say that an edge $(u,v)\in E(G)$ is *irrepairable* with respect if u or v is irrepairable, or if the corresponding edge $e'=(u',v')\in E(G')$ has $u'\notin O_{u,S}$ or $v'\notin I_{v,S}$. We let $S_{irr}\subset E(G)$ be the set of irrepairable edges after we remove S from G. We begin the proof by first proving two claims.

Claim B.1. Let P be a path of length d in $G - S_{irr}$. Then there exists a path of length *at least* d in $G' - S$.

Proof. In $G - S_{irr}$ we have removed all of the irreparable edges, so any path in the graph contains only repairable edges. By definition, if (u,v) is a repairable edge, both u and v will be repairable, and $(u, \pi_{out,u}(v)) \in O_{u,S}$ and $(v, \pi_{in,v}(u)) \in I_{v,S}$. Thus the edge corresponding to (u,v) in $G' - S$ will connect the metanodes of u and v, and $(u, \pi_{out,u}(v))$ connects to every node in $I_{u,S}$ and $(v, \pi_{out,v}(u))$ connects to every node in $O_{v,S}$. Thus the edges in $G' - S$ corresponding to the edges in P form a path of length at least d. $\square$

Claim B.2. Let $S_{irr} \subset E(G)$ be the set of irreparable edges with respect to the removed set S. Then

$$|S_{irr}| \leq 2|S|.$$

Proof. If a node v is repairable with respect to S then let $S_{irr,v}^{in} \subseteq E(G)$ (resp. $S_{irr,v}^{out}$) denote the subset of edges $(u,v) \in E(G)$ (resp. $(v,u) \in E(G)$) that are irrepairable because of S_v i.e., the corresponding edge $e' = (u', v') \in E(G')$ has $v' \notin I_{v,S}$ (resp. the corresponding edge $(v', u') \in E(G')$ has $v' \notin O_{v,S}$). Let $S_{irr,v} = S_{irr,v}^{in} \cup S_{irr,v}^{out}$. Similarly, if v is irrepairable we let $S_{irr,v} = \{(u,v) : (u,v) \in E(G)\} \cup \{(v,u) : (v,u) \in E(G)\}$ denote the set of all of v's incoming and outgoing edges. We note that $|S_{irr}| \leq \sum_v \left|S_{irr,v}\right|$ since $S_{irr} = \bigcup_v S_{irr,v}$ any irrepairable edge must be in one of the sets $S_{irr,v}$. Now we claim that $|S_{irr,v}| \leq |S_v|$ where $S_v = S \cap (\{v\} \times V_{\delta(v)})$ denote the subset of nodes deleted from the corresponding metanode. We now observe that

$$\begin{aligned} \left|S_{irr,v}\right| &\leq \left|S_{irr,v}^{in}\right| + \left|S_{irr,v}^{in}\right| \\ &\leq \left(\delta(v) - |I_{v,S}|\right) + \left(\delta(v) - |O_{v,S}|\right) \leq 2|S_v| \ . \end{aligned}$$

The last inequality invokes maximal ST-robustness to show that $\delta(v) - |O_{v,S}| \leq |S_v|$ and $\delta(v) - |I_{v,S}| \leq |S_v|$. If a node v is irrepairable then the subsets $I_{v,S}$ and $O_{v,S}$ might be empty since $\delta(v) - |S_v| \leq 0$, but it still holds that $\delta(v) - |O_{v,S}| \leq |S_v|$ and $\delta(v) - |I_{v,S}| \leq |S_v|$.

Thus

$$|S_{irr}| \leq \sum_v \left|S_{irr,v}\right| \leq \sum_v 2|S_v| \leq 2|S| \ .$$

$\square$

$\square$

Reminder of Corollary 2.5. *(of Theorem 2.3) Suppose that there exists a family $\mathbb{M} = \{M_k\}_{k=1}^{\infty}$ of max ST-robust graphs with depth d_k and constant* *indegree. Given any (e,d)-edge-depth-robust DAG G with n nodes and maximum degree δ we can construct a DAG G' with $n\times|M_\delta|$ nodes and constant indegree that is $(e/2,d\cdot d_\delta)$-depth robust.*

Proof of Corollary 2.5. (sketch) We slightly modify our reduction. Instead of replacing each node $v\in G$ with a copy of $M_{\delta(v)}$, we instead replace each node with a copy of $M_{\delta,v}:=M_\delta$, attaching the edges same way as in Construction 2.2. Thus the transformed graph G' has $|V(G)|\times|M_\delta|$ nodes and constant indegree. Let $S\subset V(G')$ be a set of nodes that we will remove from G'. By Claim B.1, there exists a path P in $G'-S$ that passes through d metanodes $M_{\delta,v_1},...,M_{\delta,v_d}$. The only difference is that each M_{δ,v_i} is maximally ST-robust *with depth d_δ* meaning we can assume that the sub-path $P_i=P\cap M_{\delta,v_i}$ through each metanode has length $|P_i|\geq d_\delta$. Thus, the total length of the path is at least $\sum_i|P_i|\geq d\cdot d_\delta$. □

Reminder of Lemma 3.3 [Sch83]. *For some suitable constant $c>0$ any any subset S of $cn/2$ vertices of G_n the graph H_n^1-S contains $k=cn^{1/3}/2$ vertex disjoint paths $A_1,...,A_k$ of length $n^{2/3}$ and H_n^2-S contains k vertex disjoint paths $B_1,...,B_k$ of the same length.*

Proof of Lemma 3.3 [Sch83]. Consider H_n^1-S. Since H_n^1 is $(cn,n^{2/3})$-depth-robust and $|S|=cn/2$, there must exist a path $A_1=(v_1,...,v_{n^{2/3}})$ in H_n^1-S. Remove all vertices of A_1 and repeat to find $A_2,...$. Then we finish with $cn/(2n^{2/3})=cn^{1/3}/2$ vertex disjoint paths of length $n^{2/3}$. We perform the same process on H_n^2 to find the B_i. □

Reminder of Lemma 3.4. *Let G_n be defined as in Construction 3.2. Then for some constants $c>0$, with high probability G_n has the property that for all $S\subset V(G_n)$ with $|S|=cn/2$ there exists $A\subseteq V(H_n^1)$ and $B\subseteq V(H_n^3)$ such that for every pair of nodes $u\in A$ and $v\in B$ the graph G_n-S contains a path from u to v and $|A|,|B|\geq 9cn/40$.*

Proof of Lemma 3.4. By Lemma 3.3, we know that in G_n-S there exists $k:=cn^{1/3}/2$ vertex disjoint paths $P_1^i,...,P_k^i$ in each H_n^i of length $n^{2/3}$. Here, c is the constant from Lemma 3.3. Let $U_{j,S}^i$ be the upper half of the j-th path in H_n^i and $L_{j,S}^i$ be the lower half, both of which are relative to the removed set S.

Let $D_{i,j,S}^{lower}$ (resp. $D_{i,j,S}^{upper}$) be an indicator random variable the event that $U_{j,S}^1$ (resp. $U_{j,S}^2$) is disconnected from $L_{i,S}^2$ (resp. $L_{i,S}^3$). Now for each $i\leq k$ define the event $BAD_{i,S}^{upper}$ to be the event that $\sum_j D_{i,j,S}^{upper}\geq k/10$ i.e., the lower path $L_{i,s}^3$ in H_n^3 is disconnected from at least $k/10$ distinct upper paths $U_{j,S}^2$ in H_n^2. Similarly, define $BAD_{j,S}^{lower}$ to be the event that $\sum_i D_{i,j,S}^{lower}\geq k/10$ i.e., the upper path $U_{j,S}^1$ is disconnected from at least $k/10$ distinct lower paths $L_{i,S}^2$ in H_n^2.

We now set $GOOD_S^{upper}=[k]\setminus\{i\ :\ BAD_{i,S}^{upper}\}$ and $GOOD_S^{lower}=[k]\setminus\{i\ :$

$BAD^{lower}_{i,S}\}$ and define

$$B_S := \bigcup_{i \in GOOD_S^{upper}}^{k} U^3_{i,S} \ , \quad \text{and} \quad A_S := \bigcup_{i \in GOOD_S^{lower}}^{k} L^1_{i,S} \ .$$

Now we claim that for every node $u \in A_S$ and $v \in B_S$ the graph $G_n - S$ contains a path from u to v. Since $u \in A_S$ we have $u \in L^1_{i,S}$ for some $i \in GOOD_S^{lower}$. Similarly, $v \in U^3_{i',S}$ for some $i' \in GOOD_S^{upper}$. By the pigeonhole principle there must exist some j s.t. $U^2_{j,S}$ connects to $L^3_{i',S}$ and $U^1_{i,S}$ connects to $L^2_{j,S}$. Thus, we can connect u to v by routing from u to $U^1_{i,S}$ to $L^2_{j,S}$ to $U^2_{j,S}$ to $L^3_{i',S}$ and finally to v. Thus, every pair of nodes in A_S and B_S are connected.

It remains to argue that (whp) for any set S the resulting set $|B_S| = |GOOD_S^{upper}| n^{2/3}$ and $|A_S| = |GOOD_S^{lower}| n^{2/3}$ are sufficiently large. Now we define the events

$$BAD_S^{lower} := |\{i \ : \ BAD^{lower}_{i,S}\}| > \frac{k}{10}$$
$$BAD_S^{upper} := |\{i \ : \ BAD^{upper}_{i,S}\}| > \frac{k}{10}.$$

Intuitively, BAD_S occurs when more than a small fraction of the events $BAD_{i,S}$ occur. Assuming that BAD_S^{upper} never occurs then for any set S we have

$$|B_S| \geq |GOOD_S| n^{2/3} \geq (9/10) k n^{2/3}/2 = 9cn/40 \ .$$

Similarly, if BAD_S^{lower} never occurs for any set S we are guaranteed to have $|A_S| \geq 9cn/40$.

Consider, for the sake of finding the probabilities, that S is fixed before all of the random edges are added to G_n. We will then union bound over all choices of sets S. First we bound $\mathbb{P}\left[BAD^{upper}_{i,S}\right]$ and $\mathbb{P}\left[BAD^{lower}_{i,S}\right]$. Union bounding over all $\binom{k}{10}$ subsets we have

$$\begin{aligned}
\mathbb{P}\left[BAD^{upper}_{i,S}\right] &\leq \binom{k}{10} \left(1 - c/40\right)^{\tau n^{2/3}/2} \\
&\leq e^k \left(\frac{1}{e}\right)^{c\tau n^{2/3}/80} \ .
\end{aligned}$$

A slightly different calculation holds for $\mathbb{P}\left[BAD^{lower}_{i,S}\right]$ since we connect each node in H^2_n to τ random nodes in H^1_n and we are now considering the upper path $U^1_{j,S}$ instead of the lower path $L^3_{i,S}$.

$$\begin{aligned}\mathbb{P}\left[BAD_{i,S}^{upper}\right] &\leq \binom{k}{10}\left(1-\frac{1}{2n^{1/3}}\right)^{\tau(k/10)n^{2/3}/2} \\ &\leq e^{k}\left(\frac{1}{e}\right)^{c\tau n^{2/3}/80}.\end{aligned}$$

By selecting $\tau > 81\cdot 80/c^2$ to ensure that $e^k\left(\frac{1}{e}\right)^{c\tau n^{2/3}/80} \leq e^{-80n^{2/3}/c}$.

We remark that for $i \neq j$ the event $BAD_{i,S}^{upper}$ is independent of $BAD_{j,S}^{upper}$ since the τ random incoming edges connected to L_i^2 are sampled independently of the edges for L_j^2.

We will show that the probability of the event BAD_S^{upper} is very small and then take a union bound over all possible S to show our desired result.

$$\begin{aligned}\mathbb{P}\left[BAD_S^{upper}\right] &\leq \binom{k}{k/10}\mathbb{P}\left[BAD_{1,S}\wedge \ldots \wedge BAD_{k/10,S}\right] \\ &= \binom{k}{k/10}\mathbb{P}\left[BAD_{1,S}^{upper}\right]^{k/10} \\ &\leq e^{k}\left[\left(\frac{1}{e}\right)^{80n^{2/3}/c}\right]^{\frac{cn^{1/3}}{20}} \\ &= \left(\frac{1}{e}\right)^{4n-k}.\end{aligned}$$

Finally, we take the union bound over every possible S of size $cn/2$ nodes. Since G_n has $2n$ nodes there are at most $2^{2n} \leq e^{2n}$ such sets. Thus,

$$\mathbb{P}\left[\exists S \text{ s.t. } BAD_S^{upper}\right] \leq e^{2n}\mathbb{P}\left[BAD_S^{upper}\right] \leq \left(\frac{1}{e}\right)^{2n-k} \ll e^{-n}.$$

Thus, except with negigible probability for any S of size $cn/2$ the event BAD_S^{upper} does not occur for any set S selected *after* G_n is sampled. Similarly, we can reason about the event BAD_S^{lower}. We now utilize the fact that

$$\mathbb{P}\left[BAD_{j,S}^{lower} : BAD_{1,S}^{lower},\ldots,BAD_{j-1,S}^{lower}\right] \leq \mathbb{P}\left[BAD_{j,S}^{lower}\right].$$

Intuitively, this holds because the event $BAD_{i,S}^{lower}$ means that for some set of $k/10$ lower paths (WLOG say $L_{1,S}^2,\ldots,L_{k/10,S}^2$) we are guaranteed that none of the edges these paths hit $U^1_{i,S}$ which only makes it more likely that those edges hit $U^1_{j,S}$ potentially preventing the event $BAD^{lower}_{j,S}$ from occuring.

$$\begin{aligned}
\mathbb{P}\Big[BAD^{lower}_S\Big] \leq& \binom{k}{k/10}\mathbb{P}\Big[BAD^{lower}_{1,S}\wedge\ldots\wedge BAD^{lower}_{k/10,S}\Big]\\
=& \binom{k}{k/10}\mathbb{P}\Big[BAD^{lower}_{1,S}\Big]\prod_{j>1}\mathbb{P}\Big[BAD^{lower}_{j,S} \ : \ BAD^{lower}_{1,S},\ldots,BAD^{lower}_{j-1,S}\Big]\\
\leq& \binom{k}{k/10}\mathbb{P}\Big[BAD^{lower}_{1,S}\Big]^{k/10}\\
\leq& e^k\left[\left(\frac{1}{e}\right)^{80n^{2/3}/c}\right]^{\frac{cn^{1/3}}{20}}\\
=& \left(\frac{1}{e}\right)^{4n-k}.
\end{aligned}$$

Thus, it follows that $\mathbb{P}\big[\exists S \text{ s.t. } BAD^{upper}_S\big] \leq e^{-n}$. $\square$