

Towards Tight(er) Bounds for the Excluded Grid Theorem*

Julia Chuzhoy[†] Zihan Tan[‡]

April 9, 2019

Abstract

We study the Excluded Grid Theorem, a fundamental structural result in graph theory, that was proved by Robertson and Seymour in their seminal work on graph minors. The theorem states that there is a function $f : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$, such that for every integer $g > 0$, every graph of treewidth at least $f(g)$ contains the $(g \times g)$ -grid as a minor. For every integer $g > 0$, let $f(g)$ be the smallest value for which the theorem holds. Establishing tight bounds on $f(g)$ is an important graph-theoretic question. Robertson and Seymour showed that $f(g) = \Omega(g^2 \log g)$ must hold. For a long time, the best known upper bounds on $f(g)$ were super-exponential in g . The first polynomial upper bound of $f(g) = O(g^{98} \text{poly log } g)$ was proved by Chekuri and Chuzhoy. It was later improved to $f(g) = O(g^{36} \text{poly log } g)$, and then to $f(g) = O(g^{19} \text{poly log } g)$. In this paper we further improve this bound to $f(g) = O(g^9 \text{poly log } g)$. We believe that our proof is significantly simpler than the proofs of the previous bounds. Moreover, while there are natural barriers that seem to prevent the previous methods from yielding tight bounds for the theorem, it seems conceivable that the techniques proposed in this paper can lead to even tighter bounds on $f(g)$.

*Extended abstract appeared in SODA 2019.

[†]Toyota Technological Institute at Chicago. Email: cjulia@ttic.edu. Part of the work was done while the author was a Weston visiting professor in the Department of Computer Science and Applied Mathematics, Weizmann Institute. Supported in part by NSF grants CCF-1318242 and CCF-1616584.

[‡]Computer Science Department, University of Chicago. Email: zihantan@uchicago.edu. Supported in part by NSF grants CCF-1318242 and CCF-1616584.

1 Introduction

The Excluded Grid theorem is a fundamental result in graph theory, that was proved by Robertson and Seymour [RS86] in their Graph Minors series. The theorem states that there is a function $f : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$, such that for every integer $g > 0$, every graph of treewidth at least $f(g)$ contains the $(g \times g)$ -grid as a minor. The theorem has found many applications in graph theory and algorithms, including routing problems [RS95], fixed-parameter tractability [DH07a, DH07b], and Erdos-Pósa-type results [RS86, Tho88, Ree97, FST11]. For an integer $g > 0$, let $f(g)$ be the smallest value, such that every graph of treewidth at least $f(g)$ contains the $(g \times g)$ -grid as a minor. An important open question is establishing tight bounds on f . Besides being a fundamental graph-theoretic question in its own right, improved upper bounds on f directly affect the running times of numerous algorithms that rely on the theorem, as well as parameters in various graph-theoretic results, such as, for example, Erdos-Pósa-type results.

On the negative side, it is easy to see that $f(g) = \Omega(g^2)$ must hold. Indeed, the complete graph on g^2 vertices has treewidth $g^2 - 1$, while the size of the largest grid minor in it is $(g \times g)$. Robertson et al. [RST94] showed a slightly stronger bound of $f(g) = \Omega(g^2 \log g)$, by using $\Omega(\log n)$ -girth constant-degree expanders, and they conjectured that this bound is tight. Demaine et al. [DHK09] conjectured that $f(g) = \Theta(g^3)$.

On the positive side, for a long time, the best known upper bounds on $f(g)$ remained super-exponential in g : the original bound of [RS86] was improved by Robertson, Seymour and Thomas in [RST94] to $f(g) = 2^{O(g^5)}$. It was further improved to $f(g) = 2^{O(g^2/\log g)}$ by Kawarabayashi and Kobayashi [KK12] and by Leaf and Seymour [LS15]. The first polynomial upper bound of $f(g) = O(g^{98} \text{poly log } g)$ was proved by Chekuri and Chuzhoy [CC16]. The proof is constructive and provides a randomized algorithm that, given an n -vertex graph G of treewidth k , finds a model of the $(g \times g)$ -grid minor in G , with $g = \tilde{\Omega}(k^{1/98})$, in time polynomial in both n and k . Unfortunately, the proof itself is quite complex. In a subsequent paper, Chuzhoy [Chu15] suggested a relatively simple framework for the proof of the theorem, that can be used to obtain a polynomial bound $f(g) = O(g^c)$ for some constant c . Using this framework, she obtained an upper bound of $f(g) = O(g^{36} \text{poly log } g)$, but unfortunately the attempts to optimize the constant in the exponent resulted in a rather technical proof. Combining the ideas from [CC16] and [Chu15], the upper bound was further improved to $f(g) = O(g^{19} \text{poly log } g)$ in [Chu16]. We note that the results in [Chu15] and [Chu16] are existential.

The main result of this paper is the proof of the following theorem.

Theorem 1.1 *There exist constants $c_1, c_2 > 0$, such that for every integer $g \geq 2$, every graph of treewidth at least $k = c_1 g^9 \log^{c_2} g$ contains the $(g \times g)$ -grid as a minor.*

Aside from significantly improving the best current upper bounds for the Excluded Grid Theorem, we believe that our framework is significantly simpler than the previous proofs. Even though a relatively simple strategy for proving the Excluded Grid Theorem was suggested in [Chu15], this strategy only led to weak polynomial bounds on $f(g)$, and obtaining tighter bounds required technically complex proofs. For example, the best previous bound of $f(g) = O(g^{19} \text{poly log } g)$ is an 80-page manuscript. Moreover, there are natural barriers that we discuss below, that prevent the strategy proposed in [Chu15] from yielding tight bounds on $f(g)$, while it is conceivable that the approach proposed in this paper will lead to even tighter bounds on $f(g)$.

Our Techniques. We now provide an overview of our techniques, and of the techniques employed in the previous proofs [CC16, Chu15, Chu16] that achieve polynomial bounds on $f(g)$.

One of the central graph-theoretic notions that we use is that of well-linkedness. Informally, we say that a subset T of vertices of a graph G is well-linked if the vertices of T are, in some sense, well-

connected in G . Formally, for every pair $T', T'' \subseteq T$ of disjoint subsets of T with $|T'| = |T''|$, there must be a collection \mathcal{P} of paths connecting every vertex of T' to a distinct vertex of T'' in G , such that the paths in \mathcal{P} are disjoint in their vertices — we call such a set \mathcal{P} of paths a set of *node-disjoint paths*. It is well known that, if T is the largest-cardinality subset of vertices of G , such that T is well-linked in G , then the treewidth of G is $\Theta(|T|)$ (see e.g. [Ree97]).

As in the proofs of [CC16, Chu15, Chu16], the main combinatorial object that we use is the *Path-of-Sets System*, that was introduced in [CC16]; a somewhat similar object (called a *grill*) was also studied by Leaf and Seymour [LS15]. A Path-of-Sets System \mathbb{P} of width w and length ℓ (see Figure 1(a)) consists of a sequence $\mathcal{C} = (C_1, \dots, C_\ell)$ of ℓ connected sub-graphs of the input graph G that we call clusters. For each cluster $C_i \subseteq G$, we are given two disjoint subsets $A_i, B_i \subseteq V(C_i)$ of its vertices of cardinality w each. We require that the vertices of $A_i \cup B_i$ are well-linked in C_i ¹. Additionally, for each $1 \leq i < \ell$, we are given a set \mathcal{P}_i of w node-disjoint paths, connecting every vertex of B_i to a distinct vertex of A_{i+1} . The paths in $\bigcup_i \mathcal{P}_i$ must be all mutually disjoint, and they cannot contain the vertices of $\bigcup_{i'=1}^\ell C_{i'}$ as inner vertices. Chekuri and Chuzhoy [CC16], strengthening a similar result of Leaf and Seymour [LS15], showed that, if a graph G contains a Path-of-Sets System of width g^2 and length g^2 , then G contains an $(\Omega(g) \times \Omega(g))$ -grid as a minor. Therefore, in order to prove Theorem 1.1, it is enough to show that a graph of treewidth $\Omega(g^9 \log^{c_2} g)$ contains a Path-of-Sets System of width and length $\Omega(g^2)$.

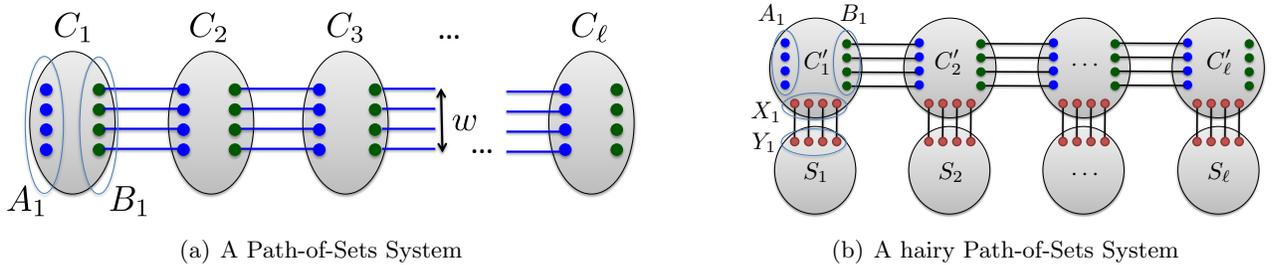


Figure 1: A Path-of-Sets System and a hairy Path-of-Sets System

Note that, if a graph G has treewidth k , then it contains a set T of $\Omega(k)$ vertices, that we call terminals, that are well-linked in G . In [CC16], the following approach was employed to construct a large Path-of-Sets System in a large-treewidth graph. Let C be any connected sub-graph of G , and let $\Gamma(C)$ be the set of the boundary vertices of C — all vertices of C that have a neighbor lying outside of C . We say that C is a *good router* if: (i) the vertices of $\Gamma(C)$ are well-linked² in C ; and (ii) there is a large set of node-disjoint paths connecting vertices of $\Gamma(C)$ to the terminals in T . The proof of [CC16] consists of two steps. First, they show that, if the treewidth of G is large, then G contains a large number of disjoint good routers. In the second step, a large subset of the good routers are combined into a Path-of-Sets System. Both these steps are quite technical, and rely on many previous results, such as the cut-matching game [KRV09], graph-reduction step preserving element-connectivity [HO96, CK14], edge-splitting [Mad78], and LP-based approximation algorithms for bounded-degree spanning tree [SL15], to name just a few. While the bound on $f(g)$ that this result produces is weak: $f(g) = O(g^{98} \text{poly } \log g)$, this result has several very useful consequences that were exploited in all subsequent proofs of the Excluded Grid Theorem, including the one in the current paper. First, the result implies that for any integer $\ell > 0$, a graph of treewidth k contains a Path-of-Sets System of length ℓ and width $\Omega(k/(\text{poly}(\ell \log k)))$. In particular, setting $\ell = \Theta(\log k)$, we

¹We use a somewhat weaker property than well-linkedness here, but for clarity of exposition we ignore these technicalities for now.

²Here, a much weaker definition of well-linkedness was used, but we ignore these technicalities in this informal overview.

can obtain a Path-of-Sets System of length ℓ and width $\Omega(k/\text{poly log } k)$. This fact was used in [CC15] to show that any graph G of treewidth k contains a sub-graph G' of treewidth $\Omega(k/\text{poly log } k)$, whose maximum vertex degree bounded by a constant, where the constant bounding the degree can be made as small as 3. This latter result proved to be a convenient starting point for subsequent improved bounds for the Excluded Grid Theorem.

In [Chu15], a different strategy for obtaining a Path-of-Sets System was suggested. Recall that, if a graph G has treewidth k , then it contains a set T of $\Omega(k)$ vertices, that we call terminals, which are well-linked in G . Partitioning the terminals into two equal-cardinality subsets A_1 and B_1 , and letting $C_1 = G$, we obtain a Path-of-Sets System of width $\Omega(k)$ and length 1. The strategy now is to perform a number of iterations, where in every iteration, the length of the current Path-of-Sets System doubles, while its width decreases by some small constant factor c . Since we eventually need to construct a Path-of-Sets System of length g^2 , we will need to perform roughly $2 \log g$ iterations, eventually obtaining a Path-of-Sets System of length g^2 and width $\Omega(k/c^{2 \log g}) = \Omega(k/g^{2 \log c})$. In order to execute a single iteration, a subroutine is employed, that, given a single cluster of the Path-of-Sets System, splits this cluster into two. Equivalently, given a Path-of-Sets System of length 1 and width w , it produces a Path-of-Sets System of length 2 and width w/c . By iteratively applying this procedure to every cluster of the current Path-of-Sets System, one obtains a new Path-of-Sets System, whose length is double the length of the original Path-of-Sets System, and the width decreases by factor c . Recall that the width of the final Path-of-Sets System that we obtain is $\Omega(k/g^{2 \log c})$, and we require that it is at least $\Omega(g^2)$, so $k \geq \Omega(g^{2 \log c+2})$ must hold. Therefore, the factor c that we lose in the splitting of a single cluster is critical for the final bound on $f(g)$, and even if this factor is quite small (which seems very non-trivial to achieve), it seems unlikely that this approach would lead to tight bounds for the Excluded Grid Theorem. The best current bound on the loss parameter c is estimated to be roughly 2^{15} . Finally, in [Chu16], the ideas from [CC16] and [Chu15] are carefully combined to obtain a tighter bound of $f(g) = \tilde{O}(g^{19})$. We note that both the results of [Chu15] and [Chu16] critically require that the maximum vertex degree of the input graph is bounded by a small constant, which can be achieved using the results of [CC15] and [CC16], as discussed above.

Our proof proceeds quite differently. Our starting point is a Path-of-Sets System of length $\ell = \Theta(\log k)$ and width $w = k/\text{poly log } k$, where k is the treewidth of the input graph G . The Path-of-Sets System can be constructed, using, e.g., the results of [CC16]. We then transform it into a structure called a *hairy* Path-of-Sets System (see Figure 1(b)) by further splitting every cluster C_i of the original Path-of-Sets System into two clusters, C'_i and S_i . The clusters C'_1, \dots, C'_ℓ are connected into a Path-of-Sets System as before, albeit with a somewhat smaller width w/c for some constant c , and for each $1 \leq i \leq \ell$, there is a set \mathcal{Q}_i of w node-disjoint paths, connecting C'_i to S_i , that are internally disjoint from both clusters. Let X_i and Y_i denote the sets of endpoints of the paths of \mathcal{Q}_i that belong to C'_i and S_i , respectively. We require that Y_i is well-linked in S_i , and that $A_i \cup B_i \cup X_i$ is well-linked in C'_i . The construction of the hairy Path-of-Sets System from the original Path-of-Sets System employs a theorem from [Chu16], that allows us to split the clusters of the Path-of-Sets System appropriately.

The main new combinatorial object that we define is a *crossbar*. Recall that we are interested in showing that the input graph G contains the $(g \times g)$ -grid as a minor. Intuitively, a crossbar inside cluster C'_i of a hairy Path-of-Sets System consists of a set \mathcal{P}_i^* of g^2 disjoint paths, connecting vertices of A_i to vertices of B_i ; and another set \mathcal{Q}_i^* of g^2 disjoint paths, where each path of \mathcal{Q}_i^* connects a distinct path of \mathcal{P}_i^* to a distinct vertex of X_i (see Figure 2). Moreover, we require that the paths of \mathcal{Q}_i^* are internally disjoint from the paths in \mathcal{P}_i^* . The main technical result that we prove is that, if the width w' of the hairy Path-of-Sets System is sufficiently large, then, for every cluster C'_i , either it contains a crossbar, or a minor of C'_i contains a Path-of-Sets System of length and width $\Omega(g^2)$. If the latter happens in any cluster C'_i , then we immediately obtain the $(g \times g)$ -grid minor inside C'_i . Therefore, we can assume that each cluster C'_i contains a crossbar. We then exploit these crossbars in

order to show that the graph G must contain an expander on $\Omega(g^2)$ vertices as a minor, such that the maximum vertex degree in the expander is bounded by $O(\log g)$. We can then employ known results to show that such an expander must contain the $(g \times g)$ -grid as a minor.

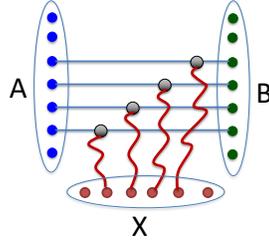


Figure 2: A Crossbar. The paths of \mathcal{P}^* are shown in blue and the paths of \mathcal{Q}^* in red.

Organization. We start with preliminaries in Section 2. In Section 3 we define a crossbar, state our main structural theorem regarding its existence, and provide the proof of Theorem 1.1 using it. We provide the proof of the structural theorem in the following two sections: in Section 4 we provide a simpler proof of a slightly weaker version of the theorem, and in Section 5 we provide its full proof.

2 Preliminaries

All logarithms in this paper are to the base of 2. All graphs are finite and they do not have loops. By default, graphs are not allowed to have parallel edges; graphs with parallel edges are explicitly called multi-graphs.

We say that a path P is *disjoint* from a set U of vertices, if $U \cap V(P) = \emptyset$. We say that it is *internally disjoint* from U , if every vertex of $U \cap V(P)$ is an endpoint of P . Given a set \mathcal{P} of paths in G , we denote by $V(\mathcal{P})$ the set of all vertices participating in the paths in \mathcal{P} . We say that two paths P, P' are *internally disjoint*, if, for every vertex $v \in V(P) \cap V(P')$, v is an endpoint of both paths. For two subsets $S, T \subseteq V(G)$ of vertices and a set \mathcal{P} of paths, we say that \mathcal{P} *connects* S to T if every path in \mathcal{P} has one endpoint in S and another in T (or it consists of a single vertex lying in $S \cap T$). We say that a set \mathcal{P} of paths is *node-disjoint* iff every pair $P, P' \in \mathcal{P}$ of distinct paths are disjoint, that is, $V(P) \cap V(P') = \emptyset$. Similarly, we say that a set \mathcal{P} of paths is *edge-disjoint* iff for every pair $P, P' \in \mathcal{P}$ of distinct paths, $E(P) \cap E(P') = \emptyset$. We sometimes refer to connected subgraphs of a given graph as *clusters*.

Treewidth, Minors and Grids. The *treewidth* of a graph $G = (V, E)$ is defined via tree-decompositions. A tree-decomposition of G consists of a tree τ , and, for each node $v \in V(\tau)$, a subset $B_v \subseteq V$ of vertices of G (called a *bag*), such that: (i) for each edge $(v, v') \in E$, there is a node $u \in V(\tau)$ with $v, v' \in B_u$; and (ii) for each vertex $v \in V$, the set $\{u \in V(\tau) \mid v \in B_u\}$ of nodes of τ induces a non-empty connected subtree of τ . The *width* of a tree-decomposition is $\max_{v \in V(\tau)} \{|B_v| - 1\}$, and the *treewidth* of a graph G , denoted by $\text{tw}(G)$, is the width of a minimum-width tree-decomposition of G .

We say that a graph H is a *minor* of a graph G , iff H can be obtained from G by a sequence of vertex deletion, edge deletion, and edge contraction operations. Equivalently, a graph H is a minor of G iff there is a function φ , mapping each vertex $v \in V(H)$ to a connected subgraph $\varphi(v) \subseteq G$, and each edge $e = (u, v) \in E(H)$ to a path $\varphi(e)$ in G connecting a vertex of $\varphi(u)$ to a vertex of $\varphi(v)$, such that: (i) for all $u, v \in V(H)$, if $u \neq v$, then $\varphi(u) \cap \varphi(v) = \emptyset$; and (ii) the paths in set $\{\varphi(e) \mid e \in E(H)\}$ are pairwise internally disjoint, and they are internally disjoint from $\bigcup_{v \in V(H)} \varphi(v)$. A map φ satisfying

these conditions is called a *model*³ of H in G . We sometimes also say that φ is an *embedding* of H into G , and, for all $v \in V$ and $e \in E$, we specifically refer to $\varphi(v)$ as the embedding of vertex v and to $\varphi(e)$ as the embedding of edge e .

The $(g \times g)$ -grid is a graph whose vertex set is: $\{v(i, j) \mid 1 \leq i, j \leq g\}$. The edge set consists of two subsets: a set of *horizontal edges* $E_1 = \{(v(i, j), v(i, j + 1)) \mid 1 \leq i \leq g; 1 \leq j < g\}$; and a set of *vertical edges* $E_2 = \{(v(i, j), v(i + 1, j)) \mid 1 \leq i < g; 1 \leq j \leq g\}$. We say that a graph G contains the $(g \times g)$ -grid minor iff some minor H of G is isomorphic to the $(g \times g)$ -grid.

Well-linkedness and Linkedness.

Definition. Let G be a graph and let T be a subset of its vertices. We say that T is *node-well-linked* in G , iff for every pair T', T'' of disjoint subsets of T , there is a set \mathcal{P} of **node-disjoint** paths in G connecting vertices of T' to vertices of T'' , with $|\mathcal{P}| = \min\{|T'|, |T''|\}$. We say that T is *edge-well-linked* in G , iff for every pair T', T'' of disjoint subsets of T , there is a set \mathcal{P}' of **edge-disjoint** paths in G connecting vertices of T' to vertices of T'' , with $|\mathcal{P}'| = \min\{|T'|, |T''|\}$. (Note that in the latter definition we allow the paths of \mathcal{P}' to share their endpoints and inner vertices).

Even though we do not use it directly, a useful fact to keep in mind is that, if T is the largest-cardinality subset of vertices of a graph G , such that T is node-well-linked, then the treewidth of G is between $|T|/4 - 1$ and $|T| - 1$ (see e.g. [Ree97]).

Definition. Let G be a graph, and let A, B be two disjoint subsets of its vertices. We say that (A, B) are *node-linked* (or simply *linked*) in G , iff for every pair $A' \subseteq A, B' \subseteq B$ of vertex subsets, there is a set \mathcal{P} of node-disjoint paths in G connecting vertices of A' to vertices of B' , with $|\mathcal{P}| = \min\{|A'|, |B'|\}$.

A Path-of-Sets System. As in previous proofs of the Excluded Grid Theorem, we rely on the notion of Path-of-Sets System, that we define next (see Figure 1(a)).

Definition. Given integers $\ell, w > 0$, a *Path-of-Sets System* \mathbb{P} of length ℓ and width w consists of the following three ingredients: (i) a sequence $\mathcal{C} = (C_1, \dots, C_\ell)$ of mutually disjoint clusters; (ii) for each $1 \leq i \leq \ell$, two disjoint subsets $A_i, B_i \subseteq V(C_i)$ of vertices of cardinality w each, and (iii) for each $1 \leq i < \ell$, a set \mathcal{P}_i of w node-disjoint paths connecting B_i to A_{i+1} , such that all paths in $\bigcup_{i=1}^{\ell-1} \mathcal{P}_i$ are node-disjoint, and they are internally disjoint from $\bigcup_{i=1}^{\ell} V(C_i)$. In other words, a path $P \in \mathcal{P}_i$ starts at a vertex of B_i , terminates at a vertex of A_{i+1} , and is otherwise disjoint from the clusters in \mathcal{C} .

We say that \mathbb{P} is a *weak Path-of-Sets System* iff for each $1 \leq i \leq \ell$, $A_i \cup B_i$ is edge-well-linked in C_i . We say that \mathbb{P} is a *strong Path-of-Sets System* iff for each $1 \leq i \leq \ell$, each of the sets A_i, B_i is node-well-linked in C_i , and (A_i, B_i) are linked in C_i .

We sometimes call the vertices of $\bigcup_i (A_i \cup B_i)$ the *nails of the Path-of-Sets System*.

Note that a Path-of-Sets System \mathbb{P} of length ℓ and width w is completely determined by $\mathcal{C}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, A_1$ and B_ℓ , so we will denote $\mathbb{P} = (\mathcal{C}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, A_1, B_\ell)$. The following theorem was proved in [CC16]; a similar theorem with slightly weaker bounds was proved in [LS15].

Theorem 2.1 *There is a constant $c \geq 1$, such that for every integer $g \geq 2$, and for every graph G , if G contains a strong Path-of-Sets System of length $\ell = g^2$ and width $w = g^2$, then it contains the $(g' \times g')$ -grid as a minor, for $g' = \lfloor g/c \rfloor$.*

³Note that this is somewhat different from the standard definition of a model, where for each edge $e \in E(H)$, $\varphi(e)$ is required to be a single edge, but it is easy to see that the two definitions are equivalent.

Stitching a Path-of-Sets System. Suppose we are given a Path-of-Sets System $\mathbb{P} = (\mathcal{C}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, A_1, B_\ell)$, with $\mathcal{C} = (C_1, \dots, C_\ell)$. Assume that for each odd-indexed cluster C_{2i-1} , we select some subsets $A'_{2i-1} \subseteq A_{2i-1}, B'_{2i-1} \subseteq B_{2i-1}$ of vertices of cardinality w' each, that have some special properties that we desire. We would like to construct a new Path-of-Sets System, whose clusters are all odd-indexed clusters of \mathcal{C} , and whose nails are $\bigcup_{i=1}^{\lceil \ell/2 \rceil} (A'_{2i-1} \cup B'_{2i-1})$. The stitching procedure allows us to do so, by exploiting the even-indexed clusters of \mathbb{P} as connectors. The proof of the following claim is straightforward and is deferred to the Appendix.

Claim 2.2 *Let $\mathbb{P} = (\mathcal{C}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, A_1, B_\ell)$ be a Path-of-Sets System of length ℓ and width w for some $\ell, w \geq 1$. Suppose we are given, for all $1 \leq i \leq \lceil \ell/2 \rceil$, subsets $A'_{2i-1} \subseteq A_{2i-1}, B'_{2i-1} \subseteq B_{2i-1}$ of vertices of cardinality w' each. Then there is a Path-of-Sets System $\hat{\mathbb{P}} = (\hat{\mathcal{C}}, \{\hat{\mathcal{P}}_i\}_{i=1}^{\lceil \ell/2 \rceil-1}, \hat{A}_1, \hat{B}_{\lceil \ell/2 \rceil})$ of length $\lceil \ell/2 \rceil$ and width w' , such that $\hat{\mathcal{C}} = (C_1, C_3, \dots, C_{2\lceil \ell/2 \rceil-1})$; and for each $1 \leq i \leq \lceil \ell/2 \rceil$, $\hat{A}_i = A'_{2i-1}$ and $\hat{B}_i = B'_{2i-1}$.*

Notice that, if \mathbb{P} is a strong Path-of-Sets System in the statement of Claim 2.2, then so is $\hat{\mathbb{P}}$.

Hairy Path-of-Sets System. Our starting point is another structure, closely related to the Path-of-Sets System, that we call a *hairy* Path-of-Sets System (see Figure 1(b)). Intuitively, the hairy Path-of-Sets System is defined similarly to a strong Path-of-Sets System, except that now, for each $1 \leq i \leq \ell$, we have an additional cluster S_i that connects to C_i with a collection of w node-disjoint paths. We require that the endpoints of these paths are suitably well-linked in C_i and S_i , respectively.

Definition. *A hairy Path-of-Sets System \mathbb{H} of length ℓ and width w consists of the following four ingredients:*

- a strong Path-of-Sets System $\mathbb{P} = (\mathcal{C}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, A_1, B_\ell)$ of length ℓ and width w ;
- a sequence $\mathcal{S} = (S_1, \dots, S_\ell)$ of disjoint clusters, such that each cluster S_i is disjoint from $\bigcup_{j=1}^{\ell} V(C_j)$ and from $\bigcup_{j=1}^{\ell-1} V(\mathcal{P}_j)$;
- for each $1 \leq i \leq \ell$, a set $Y_i \subseteq V(S_i)$ of w vertices that are node-well-linked in S_i , and a set $X_i \subseteq V(C_i)$ of w vertices, such that $X_i \cap (A_i \cup B_i) = \emptyset$, and (A_i, X_i) are node-linked in C_i ; and
- for each $1 \leq i \leq \ell$, a collection \mathcal{Q}_i of w node-disjoint paths connecting X_i to Y_i , such that all paths in $\bigcup_{j=1}^{\ell} \mathcal{Q}_j$ are disjoint from each other and from the paths in $\bigcup_{j=1}^{\ell-1} \mathcal{P}_j$, and they are internally disjoint from $\bigcup_{j=1}^{\ell} (S_j \cup C_j)$.

Note that a hairy Path-of-Sets System \mathbb{H} of length ℓ and width w is completely determined by $\mathcal{C}, \mathcal{S}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, \{\mathcal{Q}_i\}_{i=1}^{\ell}, A_1$ and B_ℓ , so we will denote $\mathbb{H} = (\mathcal{C}, \mathcal{S}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, \{\mathcal{Q}_i\}_{i=1}^{\ell}, A_1, B_\ell)$.

We note that Chekuri and Chuzhoy [CC16] showed that for all integers $\ell, w, k > 1$ with $k/\text{poly log } k = \Omega(w^{\ell^{48}})$, every graph G of treewidth at least k contains a strong Path-of-Sets System of length ℓ and width w . We prove an analogue of this result for the hairy Path-of-Sets System. The proof mostly follows from previous work and is delayed to the Appendix. We will exploit this theorem only for the setting where $\ell = \Theta(\log k)$ and $w = k/\text{poly log } k$.

Theorem 2.3 *There are constants $c, c' > 0$, such that for all integers $\ell, w, k > 1$ with $k/\log^{c'} k > cw^{\ell^{48}}$, every graph G of treewidth at least k contains a subgraph G' of maximum vertex degree 3, such that G' contains a hairy Path-of-Sets System of length ℓ and width w .*

3 Proof of the Excluded Grid Theorem

In this section, we provide a proof of Theorem 1.1, with some details delayed to Sections 4 and 5. We start by introducing the main new combinatorial object that we use, called a *crossbar*.

Definition. Let H be a graph, let A, B, X be three disjoint subsets of its vertices, and let $\rho > 0$ be an integer. An (A, B, X) -crossbar of width ρ consists of a collection \mathcal{P}^* of ρ paths, each of which connects a vertex of A to a vertex of B , and, for each path $P \in \mathcal{P}^*$, a path Q_P , connecting a vertex of P to a vertex of X , such that:

- The paths in \mathcal{P}^* are completely disjoint from each other;
- The paths in $\mathcal{Q}^* = \{Q_P \mid P \in \mathcal{P}^*\}$ are completely disjoint from each other; and
- For each pair $P \in \mathcal{P}^*$ and $Q \in \mathcal{Q}^*$ of paths, if $Q \neq Q_P$, then P and Q are disjoint; otherwise $P \cap Q$ contains a single vertex, which is an endpoint of Q_P (see Figure 2).

The following theorem is the main technical result of this paper.

Theorem 3.1 Let H be a graph and let $g \geq 2$ be an integer, such that g is an integral power of 2. Let A, B, X be three disjoint sets of vertices of H , each of cardinality $\kappa \geq 2^{22}g^9 \log g$, such that every vertex in X has degree 1 in H . Assume further that there is a set $\tilde{\mathcal{P}}$ of κ node-disjoint paths connecting vertices of A to vertices of B in H , and a set $\tilde{\mathcal{Q}}$ of κ node-disjoint paths connecting vertices of A to vertices of X in H (but the paths $P \in \tilde{\mathcal{P}}$ and $Q \in \tilde{\mathcal{Q}}$ are not necessarily disjoint). Then, either H contains an (A, B, X) -crossbar of width g^2 , or there is a minor H' of H , that contains a strong Path-of-Sets System of length $\Omega(g^2)$ and width $\Omega(g^2)$.

We defer the proof of Theorem 3.1 to Sections 4 and 5. In Section 4 we provide a simpler proof of a slightly weaker version of Theorem 3.1, where we require that $\kappa = \Omega(g^{10} \log g)$, leading to a slightly weaker bound of $f(g) = O(g^{10} \text{poly log } g)$ for Theorem 1.1. In Section 5, we provide the full proof of Theorem 3.1.

We use the following theorem to complete the proof of Theorem 1.1.

Theorem 3.2 There is a constant \tilde{c} , such that the following holds. Let G be any graph with maximum vertex degree at most 3, such that G contains a hairy Path-of-Sets System $\mathbb{H} = (\mathcal{C}, \mathcal{S}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, \{\mathcal{Q}_i\}_{i=1}^{\ell}, A_1, B_\ell)$ of length $\ell = \tilde{c} \log g$ and width $\tilde{w} \geq g^2$, for some integer $g \geq 2$ that is an integral power of 2. Assume further that for every odd integer $1 \leq i \leq \ell$, there is an (A_i, B_i, Y_i) -crossbar in graph $C_i \cup \mathcal{Q}_i$ of width g^2 . Then G contains the $(g' \times g')$ -grid as a minor, for $g' = \Omega(g / \log^{1.5} g)$.

We first complete the proof of Theorem 1.1 using Theorems 3.2 and 3.1. Let G be any graph, and let k be its treewidth. Let $g \geq 2$ be an integer, such that g is an integral power of 2, and such that for some large enough constants c'_1, c'_2 , $k / (\log k)^{c'_2} > c'_1 g^9$. We show that G contains a grid minor of size $(\Omega(g / \text{poly log } g) \times (\Omega(g / \text{poly log } g)))$.

Let $\ell = \tilde{c} \log g = O(\log k)$, where \tilde{c} is the constant from Theorem 3.2, and let $w = 2^{22}g^9 \log g = O(g^9 \log k)$. By setting the constants c'_1 and c'_2 in the bound on k appropriately, we can ensure that the conditions of Theorem 2.3, hold for ℓ, w and k . From Theorem 2.3, there is a subgraph G' of G , of maximum vertex degree 3, such that G' contains a hairy Path-of-Sets System $\mathbb{H} = (\mathcal{C}, \mathcal{S}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, \{\mathcal{Q}_i\}_{i=1}^{\ell}, A_1, B_\ell)$ of length ℓ and width w .

Let $1 \leq i \leq \ell$ be an odd integer. Consider the graph $C_i \cup \mathcal{Q}_i$ of the hairy Path-of-Sets System. Since every pair of the vertex subsets A_i, B_i, X_i are linked in C_i , there is a set $\tilde{\mathcal{P}}_i$ of w node-disjoint paths connecting vertices of A_i to vertices of B_i , and a set $\tilde{\mathcal{Q}}'_i$ of w node-disjoint paths connecting vertices

of A_i to vertices of X_i in C_i . By concatenating the paths in \tilde{Q}'_i with the paths in Q_i , we obtain a set \tilde{Q}_i of w node-disjoint paths connecting vertices of A_i to vertices of Y_i . It is also immediate to verify that every vertex in Y_i has degree 1 in $C_i \cup Q_i$. We can therefore apply Theorem 3.1 to the graph $C_i \cup Q_i$, and the sets A_i, B_i, Y_i of its vertices. If, for any odd integer $1 \leq i \leq \ell$, the outcome of Theorem 3.1 is a strong Path-of-Sets System of length $\Omega(g^2)$ and width $\Omega(g^2)$ in some minor of $C_i \cup Q_i$, then from Theorem 2.1, graph G contains a grid minor of size $(\Omega(g) \times \Omega(g))$. Therefore, we can assume from now on that for every odd integer $1 \leq i \leq \ell$, the outcome of Theorem 3.1 is an (A_i, B_i, Y_i) -crossbar in $C_i \cup Q_i$ of width g^2 . But then from Theorem 3.2, G contains the $(g' \times g')$ -grid as a minor, for $g' = \Omega(g/\log^{1.5} g)$. We conclude that there are constants c'_1, c'_2 , such that for all integers k, g with $k/\log^{c'_2} k \geq c'_1 g^9$, if a graph G has treewidth k , then it contains a grid minor of size $(g' \times g')$, where $g' = \Omega(g/\log^5 g)$. (If g is not an integral power of 2, then we round it up to the closest integral power of 2 and absorb this additional factor of 2 in the constants c'_1 and c'_2). It follows that there are constants c''_1, c''_2 , such that for all integers k, g' with $k/\log^{c''_2}(k) \geq c''_1 (g')^9$, if a graph G has treewidth k , then it contains a grid minor of size $(g' \times g')$.

One last issue is that we need to replace the poly $\log k$ in the above bound on k by poly $\log g$, as in the statement of Theorem 1.1. Assume that we are given a graph G of treewidth k , and an integer $g \geq 2$, such that $k \geq c_1 g^9 \log^{c_2} g$ holds, for large enough constants c_1, c_2 . Let $k'(g) \leq k$ be the smallest integer for which the inequality $k'(g) \geq c_1 g^9 \log^{c_2} g$ holds. Clearly, for some large enough constant c that is independent of g , $k'(g) \leq c g^{10}$, and so $\log k'(g) = \Theta(\log g)$. The treewidth of G is at least $k'(g)$, and, by choosing the constants c_1 and c_2 appropriately, we can guarantee that $k'(g)/\log^{c''_2}(k'(g)) \geq c''_1 g^9$. From the above arguments, graph G contains the $(g \times g)$ -grid as a minor. It now remains to prove Theorem 3.2.

Proof of Theorem 3.2.

We assume that ℓ is an even integer. For every odd integer $1 \leq i \leq \ell$, let $(\tilde{P}_i^*, \tilde{Q}_i^*)$ be the (A_i, B_i, Y_i) -crossbar of width g^2 in $C_i \cup Q_i$, and let $\tilde{A}_i \subseteq A_i, \tilde{B}_i \subseteq B_i$ be the sets of endpoints of the paths in \tilde{P}_i^* , lying in A_i and B_i , respectively, so that $|\tilde{A}_i| = |\tilde{B}_i| = g^2$. Using the stitching claim (Claim 2.2), we can obtain a strong Path-of-Sets System $\mathbb{P}' = (\mathcal{C}', \{\mathcal{P}'_i\}_{i=1}^{\ell/2-1}, A'_1, B'_{\ell/2})$ of length $\ell/2$ and width g^2 , where $\mathcal{C}' = (C'_1, \dots, C'_{\ell/2})$, and for $1 \leq i \leq \ell/2$, $C'_i = C_{2i-1}$. We denote, for each $1 \leq i \leq \ell/2$, $A'_i = \tilde{A}_{2i-1}$ and $B'_i = \tilde{B}_{2i-1}$. We also denote $\tilde{\mathcal{P}}_i^* = \tilde{P}_{2i-1}^*$ and $\tilde{\mathcal{Q}}_i^* = \tilde{Q}_{2i-1}^*$. Recall that for each $1 \leq i < \ell/2$, the paths in \mathcal{P}'_i connect the vertices of B'_i to the vertices of A'_{i+1} . Combined with the clusters $S_1, S_3, \dots, S_{\ell-1}$ and the corresponding path sets $Q_1, Q_3, \dots, Q_{\ell-1}$, we now obtain a hairy Path-of-Sets System of length $\ell/2$ and width g^2 . For convenience, abusing the notation, we denote $\ell/2$ by ℓ , and we denote this new hairy Path-of-Sets System by $\mathbb{H} = (\mathcal{C}, \mathcal{S}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, \{\mathcal{Q}_i\}_{i=1}^{\ell}, A_1, B_\ell)$, where for all $1 \leq i < \ell$, paths in \mathcal{P}_i connect $B_i \subseteq V(C_i)$ to $A_{i+1} \subseteq V(C_{i+1})$.

Definition. We say that a multi-graph $H = (V, E)$ is an α -expander, iff $\min_{\substack{S \subseteq V: \\ |S| \leq |V|/2}} \left\{ \frac{|E(S, \bar{S})|}{|S|} \right\} \geq \alpha$.

We use the following lemma to show that G contains a large enough expander as a minor.

Lemma 3.3 *There is a multi-graph H , with $|V(H)| = g^2$, and maximum vertex degree $O(\log g)$, such that H is a $1/2$ -expander, and it is a minor of G .*

Proof: We use the cut-matching game of Khandekar, Rao and Vazirani [KRV09], defined as follows. We are given a set V of N nodes, where N is an even integer, and two players, the cut player and the matching player. The game is played in iterations. We start with a graph H with node set V and an empty edge set. In every iteration, some edges are added to H . The game ends when H becomes a $\frac{1}{2}$ -expander. The goal of the cut player is to construct a $\frac{1}{2}$ -expander in as few iterations

as possible, whereas the goal of the matching player is to prevent the construction of the expander for as long as possible. The iterations proceed as follows. In every iteration j , the cut player chooses a partition (Z_j, Z'_j) of V with $|Z_j| = |Z'_j|$, and the matching player chooses a perfect matching M_j that matches the nodes of Z_j to the nodes of Z'_j . The edges of M_j are then added to H . Khandekar, Rao, and Vazirani [KRV09] showed that there is a strategy for the cut player that guarantees that after $O(\log^2 N)$ iterations the graph H is a $(1/2)$ -expander. Orecchia et al. [OSVV08] strengthened this result by showing that after $O(\log^2 N)$ iterations the graph H is an $\Omega(\log N)$ -expander. We use another strengthening of the result of [KRV09], due to Khandekar et al. [KKOV07].

Theorem 3.4 ([KKOV07]) *There is a constant γ_{CMG} , and an adaptive strategy for the cut player such that, no matter how the matching player plays, after $\lceil \gamma_{\text{CMG}} \log N \rceil$ iterations, multi-graph H is a $1/2$ -expander with high probability.*

We note that the strategy of the cut player is a randomized algorithm that computes a partition (Z_j, Z'_j) of V to be used in the subsequent iteration. It is an adaptive strategy, in that the partition depends on the previous partitions and on the responses of the matching player in previous iterations. Note that the maximum vertex degree in the resulting expander is exactly $\lceil \gamma_{\text{CMG}} \log N \rceil$, since the set of its edges is a union of $\lceil \gamma_{\text{CMG}} \log N \rceil$ matchings.

We assume that the constant \tilde{c} in the statement of Theorem 3.2 is at least $8\gamma_{\text{CMG}}$, so that the length ℓ of the hairy Path-of-Sets System \mathbb{H} is at least $\lceil \gamma_{\text{CMG}} \log(g^2) \rceil$. We construct a $\frac{1}{2}$ -expander H over a set U of g^2 vertices, and embed it into G , as follows.

Let \mathcal{R} be the set of g^2 node-disjoint paths, obtained by concatenating the paths of $\mathcal{P}_1^*, \mathcal{P}_1, \dots, \mathcal{P}_{\ell-1}, \mathcal{P}_\ell^*$. For each vertex $u \in U$, we let $\varphi(u)$ be any path of \mathcal{R} , that we denote by R_u , such that for $u \neq u'$, $R_u \neq R_{u'}$. Next, we will construct the set of edges of H over the course of $\lceil \gamma_{\text{CMG}} \log(g^2) \rceil$ iterations, by running the cut-matching game. Recall that in each iteration $1 \leq j \leq \lceil \gamma_{\text{CMG}} \log(g^2) \rceil$, the cut player computes a partition (Z_j, Z'_j) of U into two equal-cardinality subsets. Our goal is to compute a perfect matching M_j between Z_j and Z'_j , whose edges will be added to H . For each edge $e \in M_j$, we will also compute its embedding $\varphi(e)$ into G , such that $\varphi(e)$ is contained in $\mathcal{Q}_j^* \cup S_j$, and for all $e \neq e'$, $\varphi(e) \cap \varphi(e') = \emptyset$.

We now fix some index $1 \leq j \leq \lceil \gamma_{\text{CMG}} \log(g^2) \rceil$, and consider iteration j of the Cut-Matching game. Then for every vertex $v \in U$, there is a path $Q_j(v) \in \mathcal{Q}_j^*$, that connects a vertex of R_v to some vertex of Y_j , that we denote by $y_j(v)$. We think of $y_j(v)$ as the *representative* of vertex v for iteration j . Note that the paths of \mathcal{Q}_j^* are internally disjoint from the paths of \mathcal{R} .

Consider the partition (Z_j, Z'_j) of U computed by the cut player. This partition naturally defines a partition (\hat{Z}_j, \hat{Z}'_j) of Y_j into two equal-cardinality subsets, where for each vertex $v \in U$, if $v \in Z_j$, then $y_j(v)$ is added to \hat{Z}_j , and otherwise it is added to \hat{Z}'_j (see Figure 3). Since the vertices of Y_j are node-well-linked in S_j , there is a set \mathcal{Q}_j'' of $|U|/2$ node-disjoint paths in S_j , connecting vertices of \hat{Z}_j to vertices of \hat{Z}'_j . The set \mathcal{Q}_j'' of paths defines a matching M_j between Z_j and Z'_j : for each path $Q \in \mathcal{Q}_j''$, if $y_j(v), y_j(v')$ are the endpoints of Q , then we add the edge $e = (v, v')$ to M_j , and we let the embedding $\varphi(e)$ of this edge be the concatenation of the paths $Q_j(v), Q, Q_j(v')$. It is easy to verify that $\varphi(e)$ connects a vertex of R_v to a vertex of $R_{v'}$, and it is internally disjoint from the paths in \mathcal{R} . Moreover, for $e \neq e'$, $\varphi(e) \cap \varphi(e') = \emptyset$. From Theorem 3.4, after $\lceil \gamma_{\text{CMG}} \log(g^2) \rceil$ iterations, graph H becomes a $\frac{1}{2}$ -expander, and we obtain a model of H in G . □

Let H be the from Lemma 3.3. Recall that H is a $1/2$ -expander on g^2 vertices, and that its maximum vertex degree is bounded by $O(\log g)$.

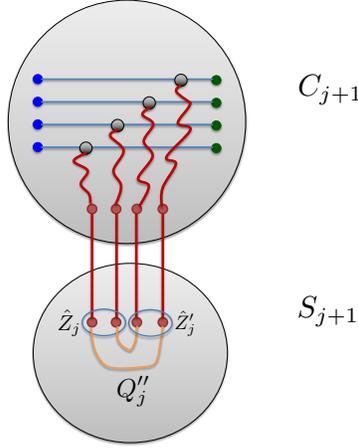


Figure 3: The execution of the j th iteration of the cut-matching game.

In general, it is well-known that a large enough expander contains a large enough grid as a minor. For example, Kleinberg and Rubinfeld [KR96] show that a bounded-degree n -vertex expander contains any graph with at most $n/\log^c n$ edges as a minor, for some constant c . Unfortunately, our expander is not bounded-degree, but has degree $O(\log n)$ (where n denotes the number of vertices in the expander). Instead, we will use a recent result of Krivelevich and Nenadov (see Theorem 8.1 in [Kri18]). We note that a similar, but somewhat weaker result was independently proved in [CN18].

Before we state the result, we need the following definition from [Kri18].

Definition. Let $G = (V, E)$ be a graph on n vertices. A vertex set $S \subseteq V$ is a separator in G if there is a partition $V = (A \cup B \cup S)$ of V , such that $|A|, |B| \leq 2n/3$, and G has no edges between A and B .

The following theorem follows from Theorem 8.1 in [Kri18] and the subsequent discussion.

Theorem 3.5 *There exist constants c, n_0 , such that the following holds. Let G be a graph on $n \geq n_0$ vertices, $0 < \alpha < 1$ a parameter, and let H be a graph with at most $cn\alpha^2/\log n$ vertices and edges. Then G has a separator of size at most cn , or G contains H as a minor.*

Let $N = g^2$ denote the number of vertices in the $1/2$ -expander H , and let $d = O(\log g)$ be the maximum vertex degree in H . We claim that every separator in H has cardinality at least $N/(24d)$. Indeed, assume otherwise, and let S be a separator of cardinality less than $N/(24d)$ in H . Let (A, B, S) be the corresponding partition of $V(H)$, so that $|A|, |B| \leq 2N/3$, and H has no edges between A and B . Assume w.l.o.g. that $|A| \leq |B|$. Since H is a $\frac{1}{2}$ -expander, at least $|A|/2$ edges are leaving A , and each such edge must have an endpoint in S . Since maximum vertex degree in H is bounded by d , $|S| \geq |A|/(2d)$ must hold. However, since we have assumed that $|S| \leq N/(24d)$, and $|B| \leq 2N/3$, we get that $|A| \geq N/6$ must hold, and so $|S| \geq N/(12d)$, a contradiction.

We can now use Theorem 3.5 in graph H , with $\alpha = 1/(24d)$, to conclude that for some constant c' , every graph H' with at most $c'g^2/\log^3 g$ edges and vertices is a minor of H . In particular, a grid of size $(g' \times g')$, where $g' = g \cdot \sqrt{c'/\log^3 g}$ is a minor of H , and hence of G . \square

4 Building the Crossbar

In this section we provide a proof of a weaker version of Theorem 3.1, where we require that $\kappa \geq 2^{22}g^{10} \log g$. We note that this weaker version of Theorem 3.1 immediately implies a slightly weaker bound of $f(g) = O(g^{10} \text{poly } \log g)$ for Theorem 1.1. Although this version of Theorem 3.1 is slightly weaker, its proof is simpler and it provides a clean framework that includes all the main ideas and the technical machinery needed for the full proof of Theorem 3.1. We complete the proof of Theorem 3.1 in Section 5.

Recall that we are given a graph H , and three disjoint subsets A, B, X of its vertices, each of cardinality $\kappa \geq 2^{22}g^{10} \log g$, such that every vertex of X has degree 1 in H . We are also given a set $\tilde{\mathcal{P}}$ of κ node-disjoint paths connecting vertices of A to vertices of B , and a set $\tilde{\mathcal{Q}}$ of κ node-disjoint paths connecting vertices of A to vertices of X . Our goal is to prove that either H contains an (A, B, X) -crossbar of width g^2 , or that its minor contains a strong Path-of-Sets system whose length and width are both at least $\Omega(g^2)$. Recall that an (A, B, X) -crossbar of width g^2 consists of a set \mathcal{P}^* of g^2 node-disjoint paths connecting vertices of A to vertices of B , and another set \mathcal{Q}^* of g^2 node-disjoint paths, such that the paths in \mathcal{Q}^* are internally disjoint from the paths in \mathcal{P}^* , and for each path $P \in \mathcal{P}^*$, there is a path $Q_P \in \mathcal{Q}^*$, whose first vertex lies on P and last vertex belongs to X .

It is easy to see that such a crossbar does not always exist. For instance, suppose H is a union of a grid of size $((\kappa + 2) \times (\kappa + 2))$ and another set X of κ vertices, where every vertex of X connects to a distinct vertex of the first column of the grid. Let A contain all vertices of the first row of the grid, excluding the grid corners, and let B contain all vertices of the last row of the grid, excluding the grid corners. The set $\tilde{\mathcal{P}}$ of paths is a subset of the columns of the grid, and the existence of the set $\tilde{\mathcal{Q}}$ of paths is easy to verify. However, there is no (A, B, X) -crossbar of width greater than 1 in this graph. We will show that in such a case, we can find a Path-of-Sets System of length and width at least $\Omega(g^2)$ in a minor of H .

Our first step is to construct two sets of paths: a set \mathcal{P} of κ node-disjoint paths connecting every vertex of A to a distinct vertex of B , and a set \mathcal{Q} of κ node-disjoint paths, connecting every vertex of A to a distinct vertex of X . Such two sets of paths are guaranteed to exist, as we can use $\mathcal{P} = \tilde{\mathcal{P}}$ and $\mathcal{Q} = \tilde{\mathcal{Q}}$. Let $H(\mathcal{P}, \mathcal{Q}) = \bigcup_{P \in \mathcal{P} \cup \mathcal{Q}} P$ be the graph obtained by the union of these paths. Among all such pairs $(\mathcal{P}, \mathcal{Q})$ of path sets, we select the sets \mathcal{P}, \mathcal{Q} that minimize the number of edges in $H(\mathcal{P}, \mathcal{Q})$. For each path $P \in \mathcal{P}$, we denote by $Q_P \in \mathcal{Q}$ the path originating at the same vertex of A as P . Even though the graph is undirected, it is convenient to think of the paths in $\mathcal{P} \cup \mathcal{Q}$ as directed away from A .

The remainder of the proof consists of six steps. In the first step, we define a new structure that we call a *pseudo-grid*. Informally, a pseudo-grid of depth D consists of a collection $\mathcal{R}_1, \dots, \mathcal{R}_D$ of disjoint subsets of paths in \mathcal{P} (that is, for all $1 \leq i \leq D$, $\mathcal{R}_i \subseteq \mathcal{P}$), such that for all i , $|\mathcal{R}_i| \leq g^2$. Additionally, if we denote $\mathcal{P}' = \mathcal{P} \setminus \bigcup_i \mathcal{R}_i$, then there must be a large subset $\mathcal{Q}' \subseteq \{Q_P \mid P \in \mathcal{P}'\}$ of paths, such that, for all $1 \leq i \leq D$, every path $Q \in \mathcal{Q}'$ intersects at least one path of \mathcal{R}_i . We show that, either H contains an (A, B, X) -crossbar of width g^2 , or it contains a pseudo-grid of a large enough depth.

In the second step, we *slice* this pseudo-grid into a large enough number M of smaller pseudo-grids. Specifically, for each path $R \in \mathcal{R}$, we define a sequence $\sigma_1(R), \dots, \sigma_M(R)$ of disjoint sub-paths of R , that appear on R in this order. Let $\Sigma_i = \{\sigma_i(R) \mid R \in \mathcal{R}\}$. For all $1 \leq i \leq M$, we let $\mathcal{Q}_i \subseteq \mathcal{Q}'$ contain only those paths Q , for which all vertices of $Q \cap V(\mathcal{R})$ belong to $V(\Sigma_i)$. We perform the slicing in a way that ensures that for all i , $|\mathcal{Q}_i|$ is large enough.

In general, for all $1 \leq i \leq M$, there are many intersections between the paths in \mathcal{Q}_i and the paths in Σ_i . But it is possible that some paths $R \in \Sigma_i$ intersect few paths of \mathcal{Q}_i and vice versa. Our third step

is a clean-up step, in which we discard all such paths, so that eventually, each path $R \in \Sigma_i$ intersects a large number of paths of \mathcal{Q}_i and vice versa.

In the fourth step, we create clusters that will be used to construct the final Path-of-Sets System. Specifically, for each $1 \leq i \leq M$, we show that there is some cluster C_i in the graph obtained from the union of the paths in Σ_i and \mathcal{Q}_i , such that there is a large enough collection $\Sigma'_i \subseteq \Sigma_i$ of paths, each of which is contained in C_i , and moreover, the endpoints of the paths in Σ'_i are well-linked in C_i . This step uses standard well-linked decomposition, though its analysis is somewhat subtle.

In the fifth step, we exploit the paths in \mathcal{R} in order to select a subset of the clusters C_i and link them into a Path-of-Sets System. Unfortunately, we will only be able to guarantee that, for each cluster C_i , the resulting vertex set $A_i \cup B_i$ is edge-well-linked in C_i ; recall that such a Path-of-Sets System is called a weak Path-of-Sets System. We then turn it into a strong Path-of-Sets System using standard techniques in our last step.

We now provide a formal proof of a weaker version of Theorem 3.1, where we assume that $\kappa \geq 2^{22}g^{10} \log g$, using the sets \mathcal{P}, \mathcal{Q} of paths that we have defined above.

4.1 Step 1: Pseudo-Grid

We define a pseudo-grid, one of our central combinatorial objects.

Definition. *Let $D > 0$ be an integer. A pseudo-grid of depth D consists of the following two ingredients. The first ingredient is a family $\{\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_D\}$ of subsets of \mathcal{P} , where for all $1 \leq i \leq D$, $|\mathcal{R}_i| \leq g^2$, and for all $1 \leq i \neq j \leq D$, $\mathcal{R}_i \cap \mathcal{R}_j = \emptyset$. Let $\mathcal{R} = \bigcup_{i=1}^D \mathcal{R}_i$, and let $\mathcal{P}' = \mathcal{P} \setminus \mathcal{R}$. The second ingredient is a set \mathcal{Q}' of $\lceil \kappa/4 \rceil$ disjoint paths, where each path $Q \in \mathcal{Q}'$ is a sub-path of a distinct path of $\{Q_P \mid P \in \mathcal{P}'\}$ (so in particular, $|\mathcal{P}'| \geq |\mathcal{Q}'| = \lceil \kappa/4 \rceil$), and exactly one endpoint of Q lies in X . Additionally, the following two properties must hold:*

P1. The paths in \mathcal{P}' are completely disjoint from the paths in \mathcal{Q}' ; and

P2. For every $1 \leq i \leq D$, the number of paths $Q \in \mathcal{Q}'$ with $Q \cap (\bigcup_{P \in \mathcal{R}_i} P) = \emptyset$ is at most $2g^2$. In other words, all but at most $2g^2$ paths of \mathcal{Q}' must intersect some path of \mathcal{R}_i .

The main result of this subsection is the following theorem.

Theorem 4.1 *Let D be any integer with $1 \leq D \leq \kappa/(2g^2)$. Then either H contains an (A, B, X) -crossbar of width g^2 , or it contains a pseudo-grid of depth D .*

We note that the theorem holds for any value of κ , and its proof does not rely on specific bounds on κ .

Proof: From the definition of the sets \mathcal{P}, \mathcal{Q} of paths, since every vertex of X has degree 1, we can assume that no path in \mathcal{P} contains a vertex of X , and that every path of \mathcal{Q} contains exactly one vertex of X , that serves as its endpoint.

We perform D iterations, where in iteration i we either construct a crossbar of width g^2 , or we compute the path set \mathcal{R}_i of the pseudo-grid. For each $1 \leq i \leq D$, we will denote by $\mathcal{P}'_i = \mathcal{P} \setminus (\mathcal{R}_1 \cup \dots \cup \mathcal{R}_i)$ the collection of the remaining paths of \mathcal{P} . We will ensure that for all i , $|\mathcal{R}_i| \leq g^2$. We let $\mathcal{P}'_0 = \mathcal{P}$.

We now describe the i th iteration of our algorithm, whose input is a set $\mathcal{P}'_{i-1} \subseteq \mathcal{P}$ of at least $\kappa - (i-1)g^2$ paths. In order to execute the i th iteration, we build a graph H_i , that is obtained from the graph H , by contracting every path $P \in \mathcal{P}'_{i-1}$ into a single vertex v_P . We keep parallel edges but discard

loops. Let $S_i = \{v_P \mid P \in \mathcal{P}'_{i-1}\}$ be the resulting set of vertices corresponding to the contracted paths. We now compute the largest set $\hat{\mathcal{Q}}$ of node-disjoint paths in H_i , connecting the vertices of S_i to the vertices of X . We consider two cases.

Case 1. The first case happens if $\hat{\mathcal{Q}}$ contains at least g^2 paths.

In this case, we show that we can construct an (A, B, X) -crossbar of width g^2 . Consider some path $Q \in \hat{\mathcal{Q}}$. We can assume without loss of generality that Q contains exactly one vertex of S_i , that serves as one of its endpoints. Let $u(Q)$ be this vertex. If $\hat{\mathcal{Q}}$ contains more than g^2 paths, we discard paths from $\hat{\mathcal{Q}}$ arbitrarily, until $|\hat{\mathcal{Q}}| = g^2$ holds. We then define \mathcal{P}^* to be the set of all paths $P \in \mathcal{P}'_{i-1}$, such that $v_P = u(Q)$ for some path $Q \in \hat{\mathcal{Q}}$. Finally, we define the set \mathcal{Q}^* of paths of the crossbar, as follows. For each path $Q \in \hat{\mathcal{Q}}$ in graph H_i , we will define a corresponding path Q' in graph H , and we will set $\mathcal{Q}^* = \{Q' \mid Q \in \hat{\mathcal{Q}}\}$. Consider now some path $Q \in \hat{\mathcal{Q}}$, and let $P \in \mathcal{P}^*$ be the path with $v_P = u(Q)$. Recall that every vertex of Q is either a vertex of H , or it is a vertex of the form $v_{P'}$ for some path $P' \in \mathcal{P}'_{i-1}$. Let $U'(Q)$ be the set of all vertices of Q that belong to H , and let $U''(Q)$ be the set of all vertices lying on the paths $P' \in \mathcal{P}'_{i-1}$, such that $v_{P'} \in V(Q)$. Finally, let $U(Q) = U'(Q) \cup U''(Q)$. Notice that for any two paths $Q, Q' \in \hat{\mathcal{Q}}$, if $Q \neq Q'$, then $U(Q) \cap U(Q') = \emptyset$, as the two paths are node-disjoint. Let $H(Q)$ be the sub-graph of H induced by the vertices of $U(Q)$. Then $H(Q)$ is a connected graph, that contains at least one vertex of P and at least one vertex of X . We let Q' be any path in $H(Q)$ connecting a vertex of P to a vertex of X , such that Q' is internally disjoint from P . Setting $\mathcal{Q}^* = \{Q' \mid Q \in \hat{\mathcal{Q}}\}$, we now obtain an (A, B, X) -crossbar $(\mathcal{P}^*, \mathcal{Q}^*)$ of width g^2 . Indeed, from the above discussion, it is immediate that the paths of \mathcal{P}^* are mutually node-disjoint, and so are the paths of \mathcal{Q}^* . From our construction, each path $Q' \in \mathcal{Q}^*$ connects a distinct path $P \in \mathcal{P}^*$ to a vertex of X . Consider now any pair $P \in \mathcal{P}^*, Q' \in \mathcal{Q}^*$ of such paths. If $v_P = u(Q')$, then from our construction $P \cap Q'$ consists of a single vertex, that serves as an endpoint of Q' . Otherwise, $Q' \cap P = \emptyset$: indeed, if $\hat{Q} \in \hat{\mathcal{Q}}$ is the path with $u(\hat{Q}) = v_P$, then, since Q' and \hat{Q} are disjoint from each other, Q' may not contain the vertex v_P , and so Q' may not contain any vertex of P .

Case 2. We now assume that $\hat{\mathcal{Q}}$ contains fewer than g^2 paths. From Menger's theorem, there is a set J_i of at most g^2 vertices in graph H_i , such that in $H_i \setminus J_i$, there is no path connecting a vertex of S_i to a vertex of X . Note that J_i may contain vertices of $S_i \cup X$.

We partition J_i into two subsets: $J'_i = J_i \cap S_i$, and $J''_i = J_i \setminus S_i$. Notice that each vertex in J''_i is also a vertex in the original graph H . We then let $\mathcal{R}_i \subseteq \mathcal{P}'_{i-1}$ be the set of all paths $P \in \mathcal{P}'_{i-1}$, whose corresponding vertex $v_P \in J'_i$. Clearly, $|\mathcal{R}_i| \leq |J_i| \leq g^2$. We define $\mathcal{P}'_i = \mathcal{P} \setminus (\mathcal{R}_1 \cup \dots \cup \mathcal{R}_i) = \mathcal{P}'_{i-1} \setminus \mathcal{R}_i$. Let $V_i = J''_i \cup (\bigcup_{P \in \mathcal{R}_i} V(P))$, a set of vertices of the original graph H , and let $\mathcal{Q}'_i = \{Q_P \mid P \in \mathcal{P}'_i\}$. Then each path in \mathcal{Q}'_i must contain a vertex of V_i . For each such path $Q \in \mathcal{Q}'_i$, let $v_i(Q)$ be the last vertex of Q that belongs to V_i , and let $\sigma_i(Q)$ be the sub-path of Q between $v_i(Q)$ and the endpoint of Q that belongs to X . Note that, as $|J''_i| \leq g^2$, for all but at most g^2 paths $Q \in \mathcal{Q}'_i$, the vertex $v_i(Q)$ lies on some path of \mathcal{R}_i . We call such a path $Q \in \mathcal{Q}'_i$ an *i-good path*. We will use the following immediate observation:

Observation 4.2 For each path $Q \in \mathcal{Q}'_i$, the segment $\sigma_i(Q)$ cannot contain any vertex of $\bigcup_{P \in \mathcal{P}'_i} V(P)$.

We continue this process for D iterations. If we did not construct an (A, B, X) -crossbar of width g^2 , then we obtain the sets $\mathcal{R}_1, \dots, \mathcal{R}_D \subseteq \mathcal{P}$ of paths, where for all i , $|\mathcal{R}_i| \leq g^2$, and we set $\mathcal{P}' = \mathcal{P}'_D$. Clearly, for all $1 \leq i \neq j \leq D$, $\mathcal{R}_i \cap \mathcal{R}_j = \emptyset$. Since $D \leq \kappa/(2g^2)$, we get that $|\mathcal{P}'| \geq \kappa/4$. We let $\mathcal{Q}' = \{\sigma_D(Q) \mid Q \in \mathcal{Q}'_D\}$. If $|\mathcal{Q}'| > \lceil \kappa/4 \rceil$, then we discard arbitrary paths from \mathcal{Q}' until $|\mathcal{Q}'| = \lceil \kappa/4 \rceil$

holds. We now claim that $\{\mathcal{R}_1, \dots, \mathcal{R}_D\}$ and \mathcal{Q}' define a pseudo-grid of depth D . Indeed, as already observed, for each $1 \leq i \leq D$, $|\mathcal{R}_i| \leq g^2$, and for all $1 \leq i \neq j \leq D$, $\mathcal{R}_i \cap \mathcal{R}_j = \emptyset$. From the definition of the segments $\sigma_D(Q)$ for $Q \in \mathcal{Q}'_D$, every path in \mathcal{Q}' has exactly one endpoint in X . From Observation 4.2, the paths in \mathcal{Q}' are disjoint from the paths in \mathcal{P}' , thus establishing Property (P1).

It now remains to establish Property (P2). Consider some index $1 \leq i \leq D$, and some path $Q \in \mathcal{Q}'$, that is both i -good and D -good. Consider the two corresponding segments $\sigma_i(Q)$ and $\sigma_D(Q)$. Recall that, since Q is i -good, $\sigma_i(Q)$ intersects some path of \mathcal{R}_i , but, since $\mathcal{R}_D \subseteq \mathcal{P}'_i$, from Observation 4.2, $\sigma_i(Q)$ cannot contain a vertex of $\bigcup_{P \in \mathcal{R}_D} V(P)$. As $\sigma_D(Q)$ is D -good, it contains a vertex of $\bigcup_{P \in \mathcal{R}_D} V(P)$. We conclude that $\sigma_i(Q) \subseteq \sigma_D(Q)$, and so $\sigma_D(Q)$ intersects some path of \mathcal{R}_i . As at most g^2 paths of \mathcal{Q}' are not i -good, and at most g^2 paths are not D -good, all but at most $2g^2$ paths of \mathcal{Q}' must intersect some path of \mathcal{R}_i . \square

We apply Theorem 4.1 to graph H with the depth parameter $D = 64g^4$. If the outcome is an (A, B, X) -crossbar of width g^2 , then we return this crossbar and terminate the algorithm. Therefore, we assume from now on that the outcome of the theorem is a pseudo-grid of depth D .

4.2 Step 2: Slicing the Paths of \mathcal{R}

Recall that for each $1 \leq i \leq D$, there are at most $2g^2$ paths $Q \in \mathcal{Q}'$, such that Q does not intersect any path of \mathcal{R}_i . We discard all such paths from \mathcal{Q}' , obtaining a set $\mathcal{Q}'' \subseteq \mathcal{Q}'$ of paths. Observe that we discard at most $2g^2D = 128g^6 < \kappa/8$ paths, and, since $\mathcal{Q}' = \lceil \kappa/4 \rceil$, we get that $|\mathcal{Q}''| \geq \kappa/8$. We now have the following property:

- I1. For each path $Q \in \mathcal{Q}''$, for every $1 \leq i \leq D$, path Q intersects at least one path of \mathcal{R}_i .

We denote $|\mathcal{R}| = N$, so $D \leq N \leq Dg^2$. Let $A' \subseteq A$ and $B' \subseteq B$ be the sets of endpoints of the paths of \mathcal{R} lying in A and B , respectively.

Let H' be the sub-graph of H , obtained by taking the union of all paths in \mathcal{R} and all paths in \mathcal{Q}'' . The next observation follows from the definition of the sets \mathcal{P} and \mathcal{Q} of paths.

Observation 4.3 *Let e be any edge of H' lying on any path in \mathcal{R} , such that e does not lie on any path Q of the original set \mathcal{Q} . Then the largest number of node-disjoint paths connecting vertices of A' to vertices of B' in $H' \setminus \{e\}$ is at most $N - 1$.*

Proof: The key is that the paths in \mathcal{Q}' are disjoint from the paths of $\mathcal{P}' = \mathcal{P} \setminus \mathcal{R}$, and so the paths in \mathcal{P}' are disjoint from the graph H' . Assume for contradiction that $H' \setminus \{e\}$ contains a set \mathcal{P}^* of N node-disjoint paths connecting vertices of A' to vertices of B' . Then we could have defined \mathcal{P} as $\mathcal{P}' \cup \mathcal{P}^*$, and the resulting graph obtained by the union of the paths in \mathcal{P} and \mathcal{Q} would have contained strictly fewer edges than the corresponding graph with the original definition of \mathcal{P} , contradicting the definition of the sets \mathcal{P}, \mathcal{Q} of paths. \square

We need the following definition.

Definition. *Given a graph \hat{H} , and two subsets Y, Z of its vertices, with $|Y| = |Z| = r$ for some integer r , we say that \hat{H} has the unique linkage property with respect to Y, Z iff there is a set \mathcal{R} of r node-disjoint paths in \hat{H} connecting every vertex of Y to a distinct vertex of Z (that we call a (Y, Z) -linkage), and moreover this set of paths is unique. We say that \hat{H} has the perfect unique linkage property with respect to Y, Z iff additionally every vertex of \hat{H} lies on some path of \mathcal{R} – the unique (Y, Z) -linkage in \hat{H} .*

Recall that graph H' is the union of the paths in \mathcal{R} and the paths in \mathcal{Q}'' . Next, we will slightly modify the graph H' by contracting some of its edges, so that the resulting graph has the perfect unique linkage property with respect to A' and B' , while preserving Property (II). We do so by performing the following two steps:

- While there is an edge $e = (u, v)$ in H' that belongs to a path of \mathcal{R} and to a path of \mathcal{Q} , contract edge e by unifying u and v ; update the corresponding paths of \mathcal{R} and \mathcal{Q} accordingly.
- While there is a vertex $u \in V(H')$ that lies on a path of \mathcal{Q}'' , but does not belong to any path in \mathcal{R} , contract any one of the (at most two) edges incident to u , by unifying u with one of its neighbors; update the corresponding path of \mathcal{Q}'' .

Let H'' be the graph obtained at the end of this procedure.

Observation 4.4 *Graph H'' is a minor of H and Property (II) holds in H'' . Moreover, H'' has the perfect unique linkage property with respect to A' and B' , with the unique linkage being \mathcal{R} .*

Proof: It is immediate to verify that H'' is a minor of H' and hence of H , and that each of the two transformation steps preserve Property (II). It is also immediate to verify that every vertex of H'' lies on some path of \mathcal{R} , as otherwise we could execute the second step.

Finally, observe that for every edge $e \in E(H'')$ that lies on a path of \mathcal{R} , e cannot lie on any path of the original set \mathcal{Q} , since each such edge was contracted by the first step. Clearly, \mathcal{R} remains an (A', B') -linkage. Assume for contradiction that a different (A', B') -linkage $\mathcal{R}' \neq \mathcal{R}$ exists in H'' . Then there is some edge e that is used by the paths in \mathcal{R} but it is not used by the paths in \mathcal{R}' . Therefore, $H'' \setminus \{e\}$ contains an (A', B') -linkage of cardinality N , and so does $H' \setminus \{e\}$. As e cannot lie on any path of the original set \mathcal{Q} , this contradicts Observation 4.3. \square

Next, we define a new combinatorial object that will be central to this step: an \hat{M} -slicing of a set of paths. Eventually, we will use this object to slice the paths of \mathcal{R} .

A convenient way to think about the \hat{M} -slicing is that it will allow us to “slice” the pseudo-grid into \hat{M} smaller such structures.

Definition. *Suppose we are given a set $\hat{\mathcal{R}}$ of node-disjoint paths, where for each path $R \in \hat{\mathcal{R}}$, one of its endpoints $a(R)$ is designated as the first endpoint of R , and the other endpoint $b(R)$ is designated as the last endpoint of R . Given an integer $\hat{M} > 0$, an \hat{M} -slicing $\Lambda = \{\Lambda(R)\}_{R \in \hat{\mathcal{R}}}$ of $\hat{\mathcal{R}}$ consists of a sequence $\Lambda(R) = (v_0(R), v_1(R), \dots, v_{\hat{M}}(R))$ of $\hat{M} + 1$ vertices of R , for every path $R \in \hat{\mathcal{R}}$, such that $v_0(R) = a(R)$, $v_{\hat{M}}(R) = b(R)$, and $v_0(R), v_1(R), \dots, v_{\hat{M}}(R)$ appear on R in this order (we allow the same vertex to appear multiple times in $\Lambda(R)$.)*

Assume now that we are given some set $\hat{\mathcal{R}}$ of node-disjoint paths, and another set $\hat{\mathcal{Q}}$ of node-disjoint paths, in some graph \hat{G} , such that each path $Q \in \hat{\mathcal{Q}}$ intersects at least one path of $\hat{\mathcal{R}}$. Assume also that we are given an \hat{M} -slicing $\Lambda = \{\Lambda(R)\}_{R \in \hat{\mathcal{R}}}$ of $\hat{\mathcal{R}}$. For all $1 \leq i \leq \hat{M}$, we denote by $\sigma_i(R)$ the sub-path of R lying strictly between $v_{i-1}(R)$ and $v_i(R)$, so it excludes these two vertices (notice that it is possible that $\sigma_i(R) = \emptyset$ if $v_{i-1}(R) = v_i(R)$, or if they are consecutive vertices on R). For each $1 \leq i \leq \hat{M}$, let $\Sigma_i = \{\sigma_i(R) \mid R \in \hat{\mathcal{R}}\}$, and let $\hat{\mathcal{Q}}_i \subseteq \hat{\mathcal{Q}}$ contain all paths $Q \in \hat{\mathcal{Q}}$ with the following property: for every path $R \in \hat{\mathcal{R}}$, for every vertex $v \in Q \cap R$, $v \in \sigma_i(R)$. Equivalently:

$$\hat{\mathcal{Q}}_i = \left\{ Q \in \hat{\mathcal{Q}} \mid (Q \cap \hat{\mathcal{R}}) \subseteq \bigcup_{\sigma \in \Sigma_i} \sigma \right\}.$$

We say that the *width* of the \hat{M} -slicing Λ with respect to $\hat{\mathcal{Q}}$ is \hat{w} iff $\min_{1 \leq i \leq \hat{M}} \{|\hat{\mathcal{Q}}_i|\} = \hat{w}$. Notice that from our definition, for all $i \neq j$, $|\hat{\mathcal{Q}}_i \cap \hat{\mathcal{Q}}_j| = \emptyset$. We now provide sufficient conditions for the existence of an \hat{M} -slicing of a given width.

Theorem 4.5 *Let \hat{G} be a graph, \hat{A}, \hat{B} two sets of its vertices of cardinality $\hat{N} > 0$ each, and assume that \hat{G} has the perfect unique linkage property with respect to (\hat{A}, \hat{B}) , with the unique linkage denoted by $\hat{\mathcal{R}}$. Assume that there is another set $\hat{\mathcal{Q}}$ of node-disjoint paths in \hat{G} , such that each path $Q \in \hat{\mathcal{Q}}$ intersects at least one path of $\hat{\mathcal{R}}$, and integers $\hat{M}, \hat{w} > 0$, such that $|\hat{\mathcal{Q}}| \geq \hat{M}\hat{w} + (\hat{M} + 1)\hat{N}$. Then there is an \hat{M} -slicing of $\hat{\mathcal{R}}$ of width \hat{w} with respect to $\hat{\mathcal{Q}}$ in \hat{G} .*

Proof: We use the following result of Robertson and Seymour (Lemma 2.5 from [RS83]); we note that the lemma appearing in [RS83] is somewhat weaker, but their proof immediately implies the stronger result that we state below; for completeness we include its proof in the Appendix.

Lemma 4.6 *Let \hat{G} be a graph, $\hat{A}, \hat{B} \subseteq V(\hat{G})$ two subsets of its vertices, such that $|\hat{A}| = |\hat{B}|$, and \hat{G} has the perfect unique linkage property with respect to (\hat{A}, \hat{B}) , with the unique (\hat{A}, \hat{B}) -linkage denoted by $\hat{\mathcal{R}}$. Then there is a bijection $\mu : V(\hat{G}) \rightarrow \{1, \dots, |V(\hat{G})|\}$ such that the following holds. For an integer $t > 0$, let S_t contain, for every path $R \in \hat{\mathcal{R}}$, the first vertex v on R with $\mu(v) \geq t$; if no such vertex exists, then we add the last vertex of R to S_t . Let $Y_t = \{v \in V(\hat{G}) \mid \mu(v) < t\}$ and $Z_t = \{v \in V(\hat{G}) \mid \mu(v) \geq t\}$. Then μ has the following properties:*

- For each path $R \in \hat{\mathcal{R}}$, for every pair v, v' of its vertices, if v' appears strictly before v on R , then $\mu(v') < \mu(v)$; and
- For every integer $0 < t \leq |V(\hat{G})|$, graph $\hat{G} \setminus S_t$ contains no path connecting a vertex of Y_t to a vertex of Z_t .

We apply Lemma 4.6 to graph \hat{G} and the sets \hat{A} and \hat{B} of its vertices, obtaining a bijection $\mu : V(\hat{G}) \rightarrow \{1, \dots, |V(\hat{G})|\}$.

Consider now an integer $1 \leq t \leq |V(\hat{G})|$, and the corresponding set S_t of vertices, that we refer to as *separator*. This separator contains exactly \hat{N} vertices – one vertex from each path $R \in \hat{\mathcal{R}}$. Recall that $\hat{G} \setminus S_t$ contains no path connecting $Y_t = \{v : \mu(v) < t\}$ and $Z_t = \{v : \mu(v) \geq t\}$.

We denote by $\mathcal{Q}^0(S_t) \subseteq \hat{\mathcal{Q}}$ the subset of all paths $Q \in \hat{\mathcal{Q}}$, such that $Q \cap S_t \neq \emptyset$, so $|\mathcal{Q}^0(S_t)| \leq \hat{N}$. Let $R \in \hat{\mathcal{R}}$ be some path with endpoints $a \in \hat{A}$, $b \in \hat{B}$, and let v be the unique vertex of R that belongs to S_t . Then v defines two sub-paths of R , as follows: $R_1(S_t)$ is the sub-path of R from a to v (including these two vertices), and $R_2(S_t)$ is similarly defined as the sub-path of R from v to b . If S_t contains b , then $R_1(S_t) = R$ and $R_2(S_t) = (b)$. Let $\mathcal{Q}^1(S_t) \subseteq \hat{\mathcal{Q}} \setminus \mathcal{Q}^0(S_t)$ be the set of all paths $Q \in \hat{\mathcal{Q}}$, such that Q intersects some path in $\{R_1(S_t) \mid R \in \hat{\mathcal{R}}\}$, and let $\mathcal{Q}^2(S_t)$ be defined similarly for $\{R_2(S_t) \mid R \in \hat{\mathcal{R}}\}$.

Notice that equivalently, $\mathcal{Q}^1(S_t)$ contains all paths $Q \in \hat{\mathcal{Q}}$, such that $Q \cap Y_t \neq \emptyset$ and $Q \cap S_t = \emptyset$. Similarly, $\mathcal{Q}^2(S_t)$ contains all paths $Q \in \hat{\mathcal{Q}}$ with $Q \cap Z_t \neq \emptyset$ and $Q \cap S_t = \emptyset$. It is easy to verify that the paths in $\mathcal{Q}^1(S_t)$ are disjoint from Z_t , and in particular $\mathcal{Q}^1(S_t) \cap \mathcal{Q}^2(S_t) = \emptyset$: otherwise, there is some path $Q \in \hat{\mathcal{Q}}$, that contains a vertex of Y_t and a vertex of Z_t , such that $Q \cap S_t = \emptyset$, contradicting the fact that Y_t and Z_t are separated in $\hat{G} \setminus S_t$. Notice that, since every path of $\hat{\mathcal{Q}}$ intersects at least one path of $\hat{\mathcal{R}}$, $(\mathcal{Q}^0(S_t), \mathcal{Q}^1(S_t), \mathcal{Q}^2(S_t))$ define a partition of $\hat{\mathcal{Q}}$.

The following observation will be useful in order to construct the \hat{M} -slicing.

Observation 4.7 *The sets $\{\mathcal{Q}^1(S_t)\}_{t \geq 1}$ of paths satisfy the following properties:*

1. $\mathcal{Q}^1(S_1) = \emptyset$, and $\mathcal{Q}^1(S_{|V(\hat{G})|})$ contains all but at most $|\hat{\mathcal{R}}|$ paths of $\hat{\mathcal{Q}}$ – the paths that intersect the vertices of \hat{B} ;
2. For all $1 \leq t < t' \leq |V(\hat{G})|$, $\mathcal{Q}^1(S_t) \subseteq \mathcal{Q}^1(S_{t'})$; and
3. For all $1 \leq t < |V(\hat{G})|$, $|\mathcal{Q}^1(S_{t+1}) \setminus \mathcal{Q}^1(S_t)| \leq 1$.

Proof: From the definition of S_1 , it contains the first vertex from each path of $\hat{\mathcal{R}}$, so $\mathcal{Q}^1(S_1) = \emptyset$. On the other hand, $S_{|V(\hat{G})|}$ contains the last vertex from each path of $\hat{\mathcal{R}}$, and, since every path of $\hat{\mathcal{Q}}$ intersects some path of $\hat{\mathcal{R}}$, $\mathcal{Q}^1(S_{|V(\hat{G})|})$ contains all paths of $\hat{\mathcal{Q}}$ except those intersecting the vertices of \hat{B} . This proves the first assertion. For the second assertion, recall that set $\mathcal{Q}^1(S_t)$ contains all paths $Q \in \hat{\mathcal{Q}}$ that intersect Y_t but are disjoint from S_t , and a path of $\mathcal{Q}^1(S_t)$ may not contain a vertex of Z_t . Each such path must intersect $Y_{t'}$, since $Y_t \subseteq Y_{t'}$, and it must be disjoint from $S_{t'}$, as $S_{t'} \subseteq S_t \cup Z_t$. Therefore, $\mathcal{Q}^1(S_t) \subseteq \mathcal{Q}^1(S_{t'})$ must hold.

We now turn to prove the third assertion. Let $Q \in \mathcal{Q}^1(S_{t+1}) \setminus \mathcal{Q}^1(S_t)$. Since $Q \notin \mathcal{Q}^1(S_t)$, either (i) $Q \cap S_t \neq \emptyset$, or (ii) $Q \cap S_t = \emptyset$ and $Q \cap Y_t = \emptyset$. If the former is true, then we say that Q is a type-1 path; otherwise we say that Q is a type-2 path. Note that $|S_t \setminus S_{t+1}| \leq 1$, from the definition of the sets $S_{t'}$, and since μ is a bijection. Since the paths in $\hat{\mathcal{Q}}$ are node-disjoint, there is at most one path of type 1 that belongs to $\mathcal{Q}^1(S_{t+1})$. We claim that no path of type 2 may belong to $\mathcal{Q}^1(S_{t+1})$.

Indeed, if $Q \in \mathcal{Q}^1(S_{t+1})$ is a type-2 path, then it must contain some vertex $v \in Y_{t+1}$. Assume that v lies on some path $R \in \hat{\mathcal{R}}$. Then $\mu(v) < t + 1$. If $\mu(v) < t$, then $v \in Y_t$, contradicting the fact that Q is of type 2. Therefore, $\mu(v) = t$. However, since Q is of type 2, $v \notin S_t$, which can only happen if some vertex v' that lies on the path R strictly before v has $\mu(v') \geq t$, and in particular $\mu(v') \geq \mu(v)$. But that is impossible from the properties of μ . \square

We now provide an algorithm to compute the \hat{M} -slicing. The algorithm performs $\hat{M} - 1$ iterations, where at the end of iteration i we produce an integer $1 \leq t_i \leq |V(\hat{G})|$, and an $(i+1)$ -slicing $\{\Lambda(R)\}_{R \in \hat{\mathcal{R}}}$ of $\hat{\mathcal{R}}$, such that the width of the slicing with respect to $\hat{\mathcal{Q}}$ is at least \hat{w} , and the following additional properties hold:

- $|\hat{\mathcal{Q}}_{i+1}| \geq |\hat{\mathcal{Q}}| - (i+2)\hat{N} - i\hat{w}$; and
- for each path $R \in \hat{\mathcal{R}}$, the vertex $v_i(R) \in \Lambda(R)$ is the unique vertex of $S_{t_i} \cap R$.

Notice that the above properties ensure that $\hat{\mathcal{Q}}_{i+1}$ contains all paths of $\mathcal{Q}^2(S_{t_i})$, except for at most \hat{N} paths, that contain the last endpoints of the paths in $\hat{\mathcal{R}}$.

Since we assumed that $|\hat{\mathcal{Q}}| \geq \hat{M}\hat{w} + (\hat{M} + 1)\hat{N}$, after $(\hat{M} - 1)$ iterations we obtain a valid \hat{M} -slicing of $\hat{\mathcal{R}}$ of width at least \hat{w} .

In order to execute the first iteration, we let $t_1 > 0$ be an integer, for which $|\mathcal{Q}^1(S_{t_1})| = \hat{w} + \hat{N}$. Such an integer must exist from Observation 4.7. For all $R \in \hat{\mathcal{R}}$, we let $v_1(R)$ be the unique vertex of R lying in $S(t_1)$, and $v_0(R), v_2(R)$ the endpoints of R lying in \hat{A} and \hat{B} , respectively. This immediately defines a 2-slicing of the paths in $\hat{\mathcal{R}}$. Recall that for each path $R \in \hat{\mathcal{R}}$, we obtain two segments: $\sigma_1(R)$, that is obtained from $R_1(S_{t_1})$ by removing its two endpoints, and $\sigma_2(R)$, obtained similarly from $R_2(S_{t_1})$. It is immediate to verify that set $\hat{\mathcal{Q}}_1$ of paths associated with this slicing contains every path $Q \in \mathcal{Q}^1(S_{t_1})$, except for those paths that contain the first endpoints of the paths in $\hat{\mathcal{R}}$. Therefore,

$|\hat{\mathcal{Q}}_1| \geq |\mathcal{Q}^1(S_{t_1})| - \hat{N} \geq \hat{w}$. Similarly, set $\hat{\mathcal{Q}}_2$ of paths associated with this slicing contains every path $Q \in \mathcal{Q}^2(S_{t_1})$, except for those paths that contain the last endpoints of the paths in $\hat{\mathcal{R}}$. Therefore, $|\hat{\mathcal{Q}}_2| \geq |\hat{\mathcal{Q}}| - |\mathcal{Q}^0(S_{t_1})| - |\mathcal{Q}^1(S_{t_1})| - |\hat{\mathcal{R}}| \geq |\hat{\mathcal{Q}}| - 3\hat{N} - \hat{w}$, as required.

We now fix some $1 \leq i < \hat{M} - 1$, and describe the $(i + 1)$ th iteration. We assume that we are given an $(i + 1)$ -slicing of $\hat{\mathcal{R}}$, with $\{v_i(R) \mid R \in \hat{\mathcal{R}}\} = S_{t_i}$, and $|\hat{\mathcal{Q}}_{i+1}| \geq |\hat{\mathcal{Q}}| - (i + 2)\hat{N} - i\hat{w} \geq \hat{N} + 2\hat{w}$.

Let t_{i+1} be the integer, for which $|\mathcal{Q}^1(S_{t_{i+1}}) \cap \hat{\mathcal{Q}}_{i+1}| = \hat{w}$. Such an integer must exist from Observation 4.7, since $|\hat{\mathcal{Q}}_{i+1}| \geq 2\hat{w} + \hat{N}$. Moreover, we are guaranteed that $t_{i+1} > t_i$, since $\hat{\mathcal{Q}}_{i+1} \subseteq \mathcal{Q}^2(S_{t_i})$.

For convenience, we denote $\mathcal{Q}^0(S_{t_i})$, $\mathcal{Q}^1(S_{t_i})$ and $\mathcal{Q}^2(S_{t_i})$ by \mathcal{Q}^0 , \mathcal{Q}^1 and \mathcal{Q}^2 , respectively.

For every path $R \in \hat{\mathcal{R}}$, vertices $v_0(R), \dots, v_i(R)$ remain the same as before. We let $v_{i+1}(R)$ be the unique vertex of $R \cap S_{t_{i+1}}$, and we let $v_{i+2}(R)$ be the endpoint of R lying in \hat{B} . We now obtain a new $(i + 2)$ -slicing $\Lambda' = \{\Lambda'(R)\}_{R \in \hat{\mathcal{R}}}$, where for each path $R \in \hat{\mathcal{R}}$, $\Lambda'(R) = (v_0(R), \dots, v_{i+1}(R))$. For convenience, let $\hat{\mathcal{Q}}'_{i+1}$ denote the original set $\hat{\mathcal{Q}}_{i+1}$, and let $\hat{\mathcal{Q}}_{i+1}$ denote the new set, defined with respect to the new slicing. Then $\hat{\mathcal{Q}}_{i+1}$ contains all paths of $\mathcal{Q}^1 \cap \hat{\mathcal{Q}}'_{i+1}$, and so, from the definition of t_{i+1} , $|\hat{\mathcal{Q}}_{i+1}| = \hat{w}$. Set $\hat{\mathcal{Q}}_{i+2}$ contains all paths of $\hat{\mathcal{Q}}'_{i+1} \setminus \hat{\mathcal{Q}}_{i+1}$, except for the paths containing the vertices of \mathcal{Q}^0 – there are at most \hat{N} such paths. Therefore, $|\hat{\mathcal{Q}}_{i+2}| \geq |\hat{\mathcal{Q}}'_{i+1}| - |\hat{\mathcal{Q}}_{i+1}| - \hat{N} \geq |\hat{\mathcal{Q}}| - (i + 2)\hat{N} - i\hat{w} - \hat{w} - \hat{N} \geq |\hat{\mathcal{Q}}| - (i + 3)\hat{N} - (i + 1)\hat{w}$, as required. \square

Let $M = 8g^4 \log g$. From Theorem 4.5, we can obtain an M -slicing $\Lambda = \{\Lambda(R)\}_{R \in \mathcal{R}}$ of \mathcal{R} , of width $w = 2^{11}g^6$ with respect to \mathcal{Q}'' , since $N \leq g^2D = 64g^6$, and so $|\mathcal{Q}''| \geq \kappa/8 \geq 2^{19}g^{10} \log g \geq M(2N + w)$. For every $1 \leq i \leq M$, we denote by $\Sigma_i = \{\sigma_i(R) \mid R \in \mathcal{R}\}$. We denote the subset $\hat{\mathcal{Q}}_i \subseteq \mathcal{Q}''$ of paths corresponding to Σ_i by \mathcal{Q}_i . We call $(\Sigma_i, \mathcal{Q}_i)$ the i th slice of Λ .

4.3 Step 3: Intersecting Path Sets

We start by defining (\hat{w}, \hat{D}) -intersecting pairs of path sets.

Definition. Let $\hat{\mathcal{R}}, \hat{\mathcal{Q}}$ be two sets of node-disjoint paths in a graph \hat{G} . Given integers $\hat{w}, \hat{D} > 0$, we say that $(\hat{\mathcal{R}}, \hat{\mathcal{Q}})$ is a (\hat{w}, \hat{D}) -intersecting pair of path sets iff each path $R \in \hat{\mathcal{R}}$ intersects at least \hat{w} distinct paths of $\hat{\mathcal{Q}}$, and each path $Q \in \hat{\mathcal{Q}}$ intersects at least \hat{D} distinct paths of $\hat{\mathcal{R}}$.

Lemma 4.8 Let $\hat{\mathcal{R}}, \hat{\mathcal{Q}}$ be two sets of node-disjoint paths in a graph \hat{G} , and let $\hat{w}, \hat{D} > 0$ be integers. Assume that each path $Q \in \hat{\mathcal{Q}}$ intersects at least $2\hat{D}$ distinct paths of $\hat{\mathcal{R}}$, and that $|\hat{\mathcal{Q}}| \geq 2|\hat{\mathcal{R}}|\hat{w}/\hat{D}$. Then there is a partition $(\hat{\mathcal{R}}', \hat{\mathcal{R}}'')$ of $\hat{\mathcal{R}}$, and a subset $\hat{\mathcal{Q}}' \subseteq \hat{\mathcal{Q}}$ of paths, such that $(\hat{\mathcal{R}}', \hat{\mathcal{Q}}')$ is a (\hat{w}, \hat{D}) -intersecting pair of path sets; $|\hat{\mathcal{Q}}'| \geq |\hat{\mathcal{Q}}|/2$; and every path in $\hat{\mathcal{R}}''$ intersects at most \hat{w} paths of $\hat{\mathcal{Q}}'$.

Proof: We start with $\hat{\mathcal{R}}' = \hat{\mathcal{R}}$ and $\hat{\mathcal{Q}}' = \hat{\mathcal{Q}}$, and then iterate, by performing one of the following two operations as long as possible:

- If there is a path $R \in \hat{\mathcal{R}}'$ intersecting fewer than \hat{w} distinct paths of $\hat{\mathcal{Q}}'$, delete R from $\hat{\mathcal{R}}'$.
- If there is a path $Q \in \hat{\mathcal{Q}}'$ intersecting fewer than \hat{D} distinct paths of $\hat{\mathcal{R}}'$, delete Q from $\hat{\mathcal{Q}}'$.

Clearly, when the algorithm terminates, $(\hat{\mathcal{R}}', \hat{\mathcal{Q}}')$ are a (\hat{w}, \hat{D}) -intersecting pair of path sets, and each path in $\hat{\mathcal{R}}'' = \hat{\mathcal{R}} \setminus \hat{\mathcal{R}}'$ intersects at most \hat{w} paths of $\hat{\mathcal{Q}}'$. It now remains to prove that $|\hat{\mathcal{Q}}'| \geq |\hat{\mathcal{Q}}|/2$.

Let $\Pi \subseteq \hat{\mathcal{R}} \times \hat{\mathcal{Q}}$ be the set of all pairs (R, Q) of paths, such that $R \cap Q \neq \emptyset$. We call each pair $(R, Q) \in \Pi$ an *intersection*. When a path R is deleted from $\hat{\mathcal{R}}'$, it participates in at most \hat{w} intersections. We say

that R is *responsible* for these intersections, and that these intersections are deleted due to R . Overall, all paths of $\hat{\mathcal{R}} \setminus \hat{\mathcal{R}}'$ may be responsible for at most $|\hat{\mathcal{R}}|\hat{w}$ intersections.

Consider now some path $Q \in \hat{\mathcal{Q}} \setminus \hat{\mathcal{Q}}'$. Originally, Q intersected at least $2\hat{D}$ paths of $\hat{\mathcal{R}}$, but at the time it was removed from $\hat{\mathcal{Q}}'$ it intersected at most \hat{D} such paths. Therefore, at least \hat{D} of its intersections were removed, and these intersections must have been removed due to paths in $\hat{\mathcal{R}} \setminus \hat{\mathcal{R}}'$. Therefore, $|\hat{\mathcal{Q}} \setminus \hat{\mathcal{Q}}'| \leq |\hat{\mathcal{R}}|\hat{w}/\hat{D} \leq \hat{\mathcal{Q}}/2$. \square

For each $1 \leq i \leq M$, We apply Lemma 4.8 to sets $(\Sigma_i, \mathcal{Q}_i)$ of paths, with parameters $\hat{w} = 4g^2$ and $\hat{D} = D/2$. Notice that $|\Sigma_i| \leq |\mathcal{R}| = N \leq Dg^2$, while $|\mathcal{Q}_i| \geq 2^{11}g^6$. It is then easy to verify that $|\mathcal{Q}_i| \geq 2|\Sigma_i|\hat{w}/\hat{D} = 16|\Sigma_i|g^2/D$. Recall that each path $Q \in \mathcal{Q}_i$ intersects at least D paths of Σ_i . Therefore, we obtain a partition (Σ'_i, Σ''_i) of Σ_i , and a subset, $\mathcal{Q}'_i \subseteq \mathcal{Q}_i$ of paths, such that $(\Sigma'_i, \mathcal{Q}'_i)$ is a $(4g^2, D/2)$ -intersecting pair of path sets, $|\mathcal{Q}'_i| \geq 2^{10}g^6$, and each path of Σ''_i intersects at most $4g^2$ paths of \mathcal{Q}'_i . For convenience, we provide the list of the main parameters of the current section in Subsection 4.7.

From now on, the remainder of the proof will only use the following facts. We are given a set \mathcal{R} of N disjoint paths, where $D \leq N \leq Dg^2$ and $D = 64g^4$. We are also given a slicing Λ of \mathcal{R} into $M = 8g^4 \log g$ slices. For each $1 \leq i \leq M$, set Σ_i contains the i th segment $\sigma_i(R)$ of every path $R \in \mathcal{R}$. We are also given a subset $\Sigma'_i \subseteq \Sigma_i$, and another set \mathcal{Q}'_i of node-disjoint paths, such that $(\Sigma'_i, \mathcal{Q}'_i)$ are $(4g^2, D/2)$ -intersecting. The paths in \mathcal{Q}'_i are disjoint from all paths in $\bigcup_{j \neq i} (\mathcal{Q}'_j \cup \Sigma_j)$. The remainder of the proof only uses these facts, and in particular it does not depend on the initial value of the parameter κ or on the cardinalities of sets \mathcal{Q}'_i of paths. We will use this fact in Section 5 when we prove the stronger version of Theorem 3.1.

4.4 Step 4: Well-Linked Decomposition

In this step, we need to use a slightly weakened definition of edge-well-linkedness, that was also used in previous work.

Definition. Let \hat{G} be a graph, T a subset of its vertices, and $\hat{w} > 0$ an integer. We say that T is \hat{w} -weakly well-linked in \hat{G} , iff for any two disjoint subsets T', T'' of T , there is a set of $\min\{|T'|, |T''|, \hat{w}\}$ edge-disjoint paths in \hat{G} , connecting vertices of T' to vertices of T'' .

The following observation is immediate.

Observation 4.9 Let \hat{G} be a graph, T a subset of its vertices, and $\hat{w} > 0$ an integer, such that $|T| \leq 2\hat{w}$. Assume further that T is \hat{w} -weakly-well-linked in \hat{G} . Then T is edge-well-linked in \hat{G} .

We will repeatedly use the following simple observation.

Observation 4.10 Let \hat{G} be a graph, T a subset of its vertices, and $\hat{w} > 0$ an integer. Assume that T is **not** \hat{w} -weakly well-linked in \hat{G} . Then there is a partition (X, Y) of $V(\hat{G})$, such that $|E(X, Y)| < \min\{\hat{w}, |T \cap X|, |T \cap Y|\}$.

Proof: Since T is not \hat{w} -weakly well-linked in \hat{G} , there are two disjoint subsets T', T'' of T , such that the largest number of edge-disjoint paths connecting vertices of T' to vertices of T'' contains fewer than $b = \min\{\hat{w}, |T'|, |T''|\}$ paths. From Menger's Theorem, there is a set E' of at most $b - 1$ edges, such that $\hat{G} \setminus E'$ contains no path connecting a vertex of T' to a vertex of T'' . Therefore, there is a partition (X, Y) of \hat{G} , with $E(X, Y) \subseteq E'$, $T' \subseteq X$, and $T'' \subseteq Y$. It is easy to verify that this partition has the required properties. \square

Let \hat{G} be a graph, and let $\hat{\Sigma}$ be a set of node-disjoint paths in \hat{G} . Given a sub-graph $C \subseteq \hat{G}$, we denote by $\hat{\Sigma}(C)$ the set of all paths $\sigma \in \hat{\Sigma}$, such that $\sigma \subseteq C$, and we denote by $\Gamma(C)$ the set of endpoints of all paths in $\hat{\Sigma}(C)$. We sometimes refer to sub-graphs $C \subseteq \hat{G}$ as *clusters*.

Given two parameters, \hat{w} and \hat{D} , we say that cluster C is *good* iff $\Gamma(C)$ is \hat{w} -weakly well-linked in C . We say that it is *happy*, if it is good, and additionally, $|\hat{\Sigma}(C)| \geq \hat{D}$. Following is the main theorem of this step.

Theorem 4.11 *Let \hat{G} be a graph, and let $\hat{\Sigma}$ and \hat{Q} be two sets of node-disjoint paths in \hat{G} (but a path in $\hat{\Sigma}$ and a path in \hat{Q} may intersect). Let $\hat{w}, \hat{D} > 0$ be integers, such that $(\hat{\Sigma}, \hat{Q})$ are $(4\hat{w}, 2\hat{D})$ -intersecting, and $\hat{D} \geq 8\hat{w}$. Then there is a collection \mathcal{C} of disjoint sub-graphs of \hat{G} , and a subset $\hat{\Sigma}' \subseteq \hat{\Sigma}$, such that:*

- *Each cluster $C \in \mathcal{C}$ is happy (that is, $|\hat{\Sigma}(C)| \geq \hat{D}$ and $\Gamma(C)$ is \hat{w} -weakly well-linked in C);*
- *$|\hat{\Sigma}'| \geq |\hat{\Sigma}|/4$; and*
- *every path $\sigma \in \hat{\Sigma}'$ belongs to a set $\hat{\Sigma}(C)$ for some $C \in \mathcal{C}$.*

Proof: Throughout the algorithm, we maintain a set \mathcal{C} of disjoint clusters of \hat{G} . Recall that $\hat{\Sigma}(C)$ contains all paths $\sigma \in \hat{\Sigma}$, such that $\sigma \subseteq C$. At the beginning, \mathcal{C} contains a single cluster \hat{G} , and $\hat{\Sigma}(\hat{G}) = \hat{\Sigma}$. We also maintain a set E' of edges that we have deleted, that is initialized to \emptyset . The algorithm is executed as long as there is some cluster $C \in \mathcal{C}$, such that $\Gamma(C)$ is not \hat{w} -weakly well-linked in C .

Let $C \in \mathcal{C}$ be any such cluster. For convenience, denote $T = \Gamma(C)$. From Observation 4.10, there is a partition (X, Y) of $V(C)$, such that $|E(X, Y)| < \min\{\hat{w}, |T \cap X|, |T \cap Y|\}$. Notice that in this case, each of $\hat{G}[X]$ and $\hat{G}[Y]$ must contain at least one path of $\hat{\Sigma}(C)$. Indeed, assume for contradiction that no path of $\hat{\Sigma}(C)$ is contained in $\hat{G}[X]$. Then for every vertex $v \in T \cap X$, path $\sigma \in \hat{\Sigma}(C)$ that contains v as an endpoint must contribute at least one edge to E' . Moreover, if both endpoints of σ belong to X , then at least two edges of σ lie in E' . Therefore, $|E'| \geq |T \cap X|$, a contradiction.

We add the edges of $E(X, Y)$ to E' . Let $J \subseteq \hat{\Sigma}(C)$ be the set of all paths that contain edges of $E(X, Y)$. Each path of $\hat{\Sigma}(C) \setminus J$ is now either contained in $\hat{G}[X]$ or in $\hat{G}[Y]$. We remove C from \mathcal{C} and replace it with $\hat{G}[X]$ and $\hat{G}[Y]$. This finishes the description of an iteration. Let \mathcal{C} be the final set of clusters at the end of the algorithm, and let $|\mathcal{C}| = r$. Then our algorithm has executed $r - 1$ iterations. Observe that in each iteration at most \hat{w} edges are added to E' , so at the end of the algorithm, $|E'| \leq (r - 1)\hat{w}$. For every cluster $C \in \mathcal{C}$, let $\text{out}(C)$ be the set of all edges of E' that are incident to C .

We partition all clusters in \mathcal{C} into two subsets: $\mathcal{C}_1 \subseteq \mathcal{C}$ contains all clusters with $|\text{out}(C)| < 4\hat{w}$, and \mathcal{C}_2 contains all remaining clusters of \mathcal{C} .

Observation 4.12 $|\mathcal{C}_1| \geq r/2$.

Proof: Assume otherwise. Then $|\mathcal{C}_2| > r/2$, and so $|E'| = \sum_{C \in \mathcal{C}} |\text{out}(C)|/2 > (r/2)(4\hat{w})/2 = r\hat{w}$. However, as observed above, $|E'| \leq (r - 1)\hat{w}$, a contradiction. \square

Observation 4.13 *Every cluster $C \in \mathcal{C}_1$ is happy.*

Proof: Consider some cluster $C \in \mathcal{C}_1$. Since the algorithm has terminated, $\Gamma(C)$ is \hat{w} -weakly-well-linked in C . It is now enough to show that $|\hat{\Sigma}(C)| \geq \hat{D}$. Recall that our procedure guarantees that

$\hat{\Sigma}(C) \neq \emptyset$. Let $\sigma \in \hat{\Sigma}(C)$ be any path. Since $(\hat{\Sigma}, \hat{Q})$ are $(4\hat{w}, 2\hat{D})$ -intersecting, there are at least $4\hat{w}$ paths in \hat{Q} that intersect σ . Since $|\text{out}(C)| < 4\hat{w}$, and the paths in \hat{Q} are disjoint, at least one such path Q is contained in C . Path Q , in turn, intersects at least $2\hat{D}$ paths of $\hat{\Sigma}$. Since $|\text{out}(C)| < 4\hat{w}$, all but at most $4\hat{w}$ these paths are contained in C . Therefore, $|\hat{\Sigma}(C)| \geq 2\hat{D} - 4\hat{w} \geq \hat{D}$, since we have assumed that $\hat{D} \geq 8\hat{w}$. \square

Corollary 4.14 $r \leq 2|\hat{\Sigma}|/\hat{D}$.

Proof: Note that from Observations 4.12 and 4.13, $\sum_{C \in \mathcal{C}_1} |\hat{\Sigma}(C)| \geq r\hat{D}/2$. On the other hand, $\sum_{C \in \mathcal{C}_1} |\hat{\Sigma}(C)| \leq |\hat{\Sigma}|$. The corollary now follows. \square

We say that a path $\sigma \in \hat{\Sigma}$ is destroyed if at least one of its edges belongs to E' ; otherwise, we say that it survives. From the above corollary, the number of paths that are destroyed is bounded by $r\hat{w} \leq 2\hat{w}|\hat{\Sigma}|/\hat{D} \leq |\hat{\Sigma}|/2$, since we have assumed that $\hat{D} \geq 8\hat{w}$. Each of the surviving paths belongs to some set $\hat{\Sigma}(C)$ for $C \in \mathcal{C}$. At least half the clusters in \mathcal{C} are happy. For a happy cluster C , $|\hat{\Sigma}(C)| \geq \hat{D}$, and for an unhappy cluster C , $|\hat{\Sigma}(C)| < \hat{D}$. We denote by $\hat{\Sigma}' \subseteq \hat{\Sigma}$ the set of all paths σ , such that $\sigma \in \hat{\Sigma}(C)$ for a happy cluster C . Then $\hat{\Sigma}'$ contains at least half of the surviving paths, and altogether, $|\hat{\Sigma}'| \geq |\hat{\Sigma}|/4$. \square

Recall that we have $M = 8g^4 \log g$ slices $\{(\Sigma_i, Q_i)\}_{1 \leq i \leq M}$. Recall that for each $1 \leq i \leq M$, we have computed subsets $\Sigma'_i \subseteq \Sigma_i$ and $Q'_i \subseteq Q_i$, such that (Σ'_i, Q'_i) are $(4g^2, D/2)$ -intersecting. Let H'_i be a graph obtained from the union of the paths in Σ_i and Q_i . Denote $\hat{D} = D/4$ and $\hat{w} = g^2$, so that (Σ'_i, Q'_i) are $(4\hat{w}, 2\hat{D})$ -intersecting. Note that $\hat{D} \geq 8\hat{w}$, since $D = 64g^4$. We apply Theorem 4.11 to graph H'_i , with $\hat{\Sigma} = \Sigma'_i$, $\hat{Q} = Q'_i$, and parameters \hat{D} , \hat{w} . Let \mathcal{C}_i be the resulting collection of happy clusters, and let $C_i \in \mathcal{C}_i$ be any such cluster. We denote by $\tilde{\Sigma}_i = \hat{\Sigma}(C_i)$. Recall that $|\tilde{\Sigma}_i| \geq D/4$, and the endpoints of the paths in $\tilde{\Sigma}_i$ are g^2 -weakly well-linked in C_i . Finally, we let $\tilde{\mathcal{C}} = \{C_1, \dots, C_M\}$. To summarize, we have obtained a collection of M clusters, one cluster per slice. Each cluster C_i contains a set $\tilde{\Sigma}_i$ of at least $D/4$ segments, whose endpoints are g^2 -weakly well-linked in C_i .

4.5 Step 5: Constructing a Weak Path-of-Sets System

In this section we construct a weak Path-of-Sets System of width g^2 and length g^2 in graph H'' . Recall that H'' is a minor of the graph H , and it is the union of the paths in \mathcal{R} and \mathcal{Q}'' . In particular, the maximum vertex degree in H'' is at most 4. We will use this fact in the last step, in order to turn the Path-of-Sets System into a strong one. For an integer $\hat{N} > 0$, we denote $[\hat{N}] = \{1, \dots, \hat{N}\}$. Let $\pi : \mathcal{R} \rightarrow [N]$ be an arbitrary bijection, mapping every path in \mathcal{R} to a distinct integer of $[N]$ (recall that $|\mathcal{R}| = N$). Following is the main theorem for this step.

Theorem 4.15 *Let $\hat{D}, \hat{N}, \hat{M}, \hat{w}$ be non-negative integers, such that (i) $\hat{N} \geq 3\hat{w}$; (ii) $\hat{D}^2 \geq 4\hat{N}\hat{w}$; and (iii) $\hat{M}\hat{D} \geq 2\hat{N}\hat{w}$ hold, and let $S_1, \dots, S_{\hat{M}}$ be subsets of $[\hat{N}]$, where for each $1 \leq i \leq \hat{M}$, $|S_i| \geq \hat{D}$. Then there are \hat{w} indices $1 \leq i_1 < i_2 < \dots < i_{\hat{w}} \leq \hat{M}$, such that for all $1 \leq j < \hat{w}$, $|S_{i_j} \cap S_{i_{j+1}}| \geq \hat{w}$.*

We prove this theorem below, after we show how to construct a weak Path-of-Sets System of length and width g^2 in H'' using it. Denote $\hat{M} = M = 8g^4 \log g$, $\hat{D} = D/4$, $\hat{N} = N$, and $\hat{w} = g^2$. Recall that we are given a set $\mathcal{C} = \{C_1, C_2, \dots, C_{\hat{M}}\}$ of clusters, where cluster C_i corresponds to slice i . Recall also that for each i , $|\tilde{\Sigma}(C_i)| \geq D/4$. We now build the subsets $S_1, \dots, S_{\hat{M}}$ of $[\hat{N}]$ as follows. Fix some $1 \leq i \leq \hat{M}$. For every path $R \in \mathcal{R}$, such that $\sigma_i(R) \in \tilde{\Sigma}(C_i)$, we add $\pi(R)$ to S_i . Notice that for all i , $|S_i| \geq D/4 = \hat{D}$. We now verify that the conditions of Theorem 4.15 hold for the chosen parameters. The first condition, $\hat{N} \geq 3\hat{w}$, is immediate from the fact that $N \geq D = 64g^4$. The second condition,

$\hat{D}^2 \geq 4\hat{N}\hat{w}$, is equivalent to: $D^2/16 \geq 4Ng^2$. Since $N \leq Dg^2$, it is enough to show that $D \geq 64g^4$, which holds from the definition of D . The third condition, $\hat{D}\hat{M} \geq 2\hat{N}\hat{w}$ is equivalent to: $DM \geq 8Ng^2$. Using the fact that $N \leq Dg^2$, and that $M \geq 8g^4$, the inequality clearly holds.

Therefore, we can now apply Theorem 4.15 to conclude that there are indices $1 \leq i_1 < i_2 < \dots < i_{g^2} \leq \hat{M}$, such that for all $1 \leq j < g^2$, $|S_{i_j} \cap S_{i_{j+1}}| \geq g^2$.

Next, we define, for all $1 \leq j \leq g^2$, subsets $T_j \subseteq S_{i_j}$ of g^2 indices, as follows. Set T_1 is an arbitrary subset of g^2 indices of S_{i_1} . For each $1 < j \leq g^2$, set T_j is an arbitrary subset of g^2 indices in $S_{i_{j-1}} \cap S_{i_j}$. For convenience, we also define a set $T_{g^2+1} = T_{g^2}$.

The clusters C'_1, \dots, C'_{g^2} of the Path-of-Sets system are defined as follows. For $1 \leq j \leq g^2$, cluster $C'_j = C_{i_j}$. Observe that for all $1 \leq j \leq g^2 + 1$, set T_j of indices defines a subset $\tilde{\mathcal{R}}^j \subseteq \mathcal{R}$ of paths: these are all paths $R \in \mathcal{R}$ with $\pi(R) \in T_j$. Therefore, $|\tilde{\mathcal{R}}^j| = g^2$, and for each path $R \in \tilde{\mathcal{R}}^j$, its segment $\sigma_{i_j}(R) \in \tilde{\Sigma}(C_{i_j})$. Moreover, if $j > 1$, then for every path $R \in \tilde{\mathcal{R}}^j$, $\sigma_{i_{j-1}}(R) \in \tilde{\Sigma}(C_{i_{j-1}})$.

Consider now some index $1 \leq j \leq g^2$. Let $\Sigma_{i_j}^A \subseteq \tilde{\Sigma}_{i_j}$ denote all segments $\sigma_{i_j}(R)$ of paths $R \in \tilde{\mathcal{R}}^j$, and let A_j be the set of vertices containing the first endpoint of each such segment. Let $\Sigma_{i_j}^B \subseteq \tilde{\Sigma}_{i_j}$ denote all segments $\sigma_{i_j}(R)$ of paths $R \in \tilde{\mathcal{R}}^{j+1}$, and let B_j be the set of vertices containing the last endpoint of each such segment. Then from Theorem 4.11, the endpoints of the paths in $\Sigma_{i_j}^A \cup \Sigma_{i_j}^B$ are g^2 -weakly well-linked in $C'_j = C_{i_j}$. Therefore, $A_j \cup B_j$ is g^2 -weakly well-linked in C'_j . Moreover, since $|A_j| = |B_j| = g^2$, from Observation 4.9, $A_j \cup B_j$ is edge-well-linked in C'_j .

It now remains to construct, for each $1 \leq j < g^2$, the set \mathcal{P}_j of disjoint paths, connecting every vertex of B_j to a distinct vertex of A_{j+1} . Recall that $|B_j| = |A_{j+1}| = |\tilde{\mathcal{R}}^{j+1}|$, and every path $R \in \tilde{\mathcal{R}}^{j+1}$ contains a single vertex b_R of B_j and a single vertex a_R of A_{j+1} . For each path $R \in \tilde{\mathcal{R}}^{j+1}$, we add the sub-paths of R between b_R and a_R to \mathcal{P}_j . It is easy to verify that all paths in set $\bigcup_j \mathcal{P}_j$ are disjoint from each other and are internally disjoint from the clusters C'_j . This is because for each $1 \leq j < g^2$, for each path $P \in \mathcal{P}_j$, P is a sub-path of some path $R \in \tilde{\mathcal{R}}^{j+1}$ spanning its segments $\sigma_{i_{j+1}}(R), \dots, \sigma_{i_{j+1}-1}(R)$, and two additional edges, one immediately preceding $\sigma_{i_{j+1}}(R)$, and one immediately following $\sigma_{i_{j+1}-1}(R)$. Since $i_1 < i_2 < \dots < i_{g^2}$, the paths in $\bigcup_j \mathcal{P}_j$ are disjoint from each other, and are internally disjoint from $\bigcup_{j'} C'_{j'}$.

In order to complete the proof of Theorem 3.1, it is now enough to prove Theorem 4.15.

Proof of Theorem 4.15. The proof consists of three steps. First, we use the sets $S_1, \dots, S_{\hat{M}}$ of indices to define a directed acyclic graph. Then we show that the size of the maximum independent set in this graph is small. We use this fact to conclude that the graph must contain a long directed path, which is then used to construct the desired sequence i_1, \dots, i_w of indices.

We start by defining a directed graph $\hat{G} = (\hat{V}, \hat{E})$, where $\hat{V} = \{1, 2, \dots, \hat{M}\}$, and for every pair $1 \leq i < j \leq \hat{M}$ of its vertices, we add a directed edge (i, j) to \hat{E} if and only if $|S_i \cap S_j| \geq \hat{w}$. It is easy to verify that \hat{G} is indeed a directed acyclic graph. We will use the following claim to bound the size of the maximum independent set in \hat{G} .

Claim 4.16 *Let $\mathcal{S} \subseteq \{S_1, \dots, S_{\hat{M}}\}$ be any collection of $r = \left\lceil 2\hat{N}/\hat{D} \right\rceil$ sets. Then there are two distinct sets $S_i, S_j \in \mathcal{S}$ with $|S_i \cap S_j| \geq \hat{w}$.*

The claim immediately implies the following corollary.

Corollary 4.17 *Let $V' \subseteq \hat{V}$ be any subset of vertices of \hat{G} , such that no two vertices in V' are connected by an edge. Then $|V'| < 2\hat{N}/\hat{D}$.*

We now turn to prove Claim 4.16.

Proof of Claim 4.16. Assume for contradiction that the claim is false. Then there must exist r subsets U_1, \dots, U_r of $[\hat{N}]$, each of which has cardinality at least \hat{D} , such that for every pair U_i, U_j of the subsets, with $i \neq j$, $|U_i \cap U_j| < \hat{w}$. We can assume without loss of generality that for each $1 \leq i \leq r$, $|U_i| = \hat{D}$. Notice that, if $\hat{D} \geq 2\hat{N}/3$, then every pair U_i, U_j of sets must share at least $\hat{N}/3 \geq \hat{w}$ elements, leading to a contradiction. Therefore, we assume from now on that $\hat{D} < 2\hat{N}/3$. Let $s = \sum_{i=1}^r |U_i| = r\hat{D}$.

For every element $j \in [\hat{N}]$, let n_j denote the total number of sets U_i with $j \in U_i$, so $\sum_{j \in [\hat{N}]} n_j = \sum_{i=1}^r |U_i| = s$. Let Π be the set of all triples (i, i', j) , with $1 \leq i, i' \leq r$ and $j \in [\hat{N}]$, such that element j belongs to both U_i and $U_{i'}$. Notice that each element j contributes exactly n_j^2 triples of the form (i, i', j) to Π , and so $|\Pi| = \sum_{j=1}^{\hat{N}} n_j^2$. Since the function $f(x) = x^2$ is convex, $\frac{1}{\hat{N}} \sum_{j=1}^{\hat{N}} n_j^2 \geq \left(\sum_{j=1}^{\hat{N}} \frac{n_j}{\hat{N}}\right)^2 = \left(\frac{s}{\hat{N}}\right)^2$. Therefore:

$$|\Pi| \geq \frac{s^2}{\hat{N}}. \quad (1)$$

On the other hand, consider any pair $1 \leq i, i' \leq r$ of indices. If $i = i'$, then U_i contributes $|U_i| = \hat{D}$ triples of the form (i, i, j) to Π , and the total number of triples of the form (i, i, j) for all $1 \leq i \leq r$ and $j \in [\hat{N}]$ is $r\hat{D} = s$. Otherwise, $i \neq i'$, and U_i and $U_{i'}$ contribute fewer than $2\hat{w}$ triples of the form (i, i', j) or (i', i, j) to Π . Therefore,

$$\begin{aligned} |\Pi| &< s + \binom{r}{2} \cdot 2\hat{w} \\ &\leq s + r^2\hat{w} \\ &= s + \hat{w} \cdot \left(\left[\frac{2\hat{N}}{\hat{D}}\right]\right)^2 \\ &\leq s + \hat{w} \cdot \frac{8\hat{N}^2}{\hat{D}^2}. \end{aligned} \quad (2)$$

(we have used the fact that $\hat{D} < 2\hat{N}/3$ for the last inequality). Altogether, we get that:

$$\frac{s^2}{\hat{N}} \leq |\Pi| < s + \frac{8\hat{w}\hat{N}^2}{\hat{D}^2}. \quad (3)$$

We claim that:

$$\frac{s^2}{\hat{N}} \geq 2s \quad (4)$$

and that:

$$\frac{s^2}{\hat{N}} \geq \frac{16\hat{w}\hat{N}^2}{\hat{D}^2} \quad (5)$$

Notice that, if both inequalities are true, then their average contradicts Inequality (3). We now show that both inequalities are indeed true. The first inequality is equivalent to $s \geq 2N$. Since $s = r\hat{D} = \left\lceil \frac{2\hat{N}}{\hat{D}} \right\rceil \cdot \hat{D} \geq 2\hat{N}$, this inequality is clearly true.

For the second inequality, since $s \geq 2\hat{N}$, we get that $\frac{s^2}{N} \geq 4\hat{N}$. Since we have assumed that $\hat{D}^2 \geq 4\hat{N}\hat{w}$, we get that $\frac{s^2}{N} \geq 4\hat{N} \geq \frac{16\hat{w}\hat{N}^2}{\hat{D}^2}$.

□

Next, we show that graph \hat{G} contains a long directed path. We say that a subset V' of vertices of a directed graph is an *independent set* iff no pair of vertices in V' is connected by an edge.

Claim 4.18 *Let $G = (V, E)$ be any directed acyclic graph on \hat{M} vertices. Let $\ell(G)$ be the length of the longest directed path in G , and let $\alpha(G)$ be the cardinality of the largest independent set in G . Then $\ell(G) \geq \hat{M}/\alpha(G)$.*

Proof: We construct a partition of the vertices of G into subsets L_1, L_2, \dots, L_h for some integer h , such that:

- For all $1 \leq i \leq h$, set L_i is non-empty and it is an independent set; and
- For all $i > 1$, for every vertex $v \in L_i$, there is some vertex $v' \in L_{i-1}$, such that the edge (v', v) belongs to G .

Notice that for all i , $|L_i| \leq \alpha(G)$, and, since every vertex of G belongs to some subset L_i , we get that $h \geq \hat{M}/\alpha(G)$. It is now immediate to construct a directed path in G of length h , by starting at any vertex $v \in L_h$, and then iteratively moving to a vertex in the previous subset that connects to the current vertex with an edge.

We now describe the construction of the sets L_i . Set L_1 contains all vertices of G with no incoming edges. It is easy to verify that it is a non-empty independent set. Assume now that we have defined the sets L_1, \dots, L_{i-1} . If every vertex of G belongs to one of these sets, then we terminate the algorithm. Otherwise, consider the graph $G_i = G \setminus (L_1 \cup \dots \cup L_{i-1})$. We then let L_i contain all vertices of G_i that have no incoming edges in G_i . Notice that L_i must be an independent set in the original graph G . Moreover, for every vertex $v \in L_i$, there must be some vertex $v' \in L_{i-1}$ with $(v', v) \in E(G)$, as otherwise v should have been added to L_{i-1} . □

We conclude that graph \hat{G} has a directed path of length at least $\hat{M}\hat{D}/2\hat{N} \geq \hat{w}$ (we have used the assumption that $\hat{M}\hat{D} \geq 2\hat{N}\hat{w}$ from the statement of Theorem 4.15). This directed path immediately defines the desired sequence $i_1, \dots, i_{\hat{w}}$ of indices, completing the proof of Theorem 4.15. □

4.6 Step 6: a Strong Path-of-Sets System

Recall that in Step 5, we have constructed a weak Path-of-Sets System of length g^2 and width g^2 , that we denote by $\mathbb{P} = (\mathcal{C}, \{\mathcal{P}_i\}_{i=1}^{g^2-1}, A_1, B_{g^2})$, in a minor H'' of H , whose maximum vertex degree is bounded by 4. Let $\ell = w = g^2$. Abusing the notation, we denote $\mathcal{C} = (C_1, \dots, C_\ell)$. In this step we complete the proof of (the weaker version of) Theorem 3.1, by converting \mathbb{P} into a strong Path-of-Sets System. The length of the new Path-of-Sets System will remain ℓ , and the set \mathcal{C} of clusters will remain the same. The width will decrease by a constant factor.

This step uses standard techniques, and is mostly identical to similar steps in previous proofs of the Excluded Grid Theorem of [CC16, Chu15, Chu16]. In particular, we will use the Boosting Theorems

of [CC16] in order to select large subsets $\mathcal{P}'_i \subseteq \mathcal{P}_i$ of paths, such that their endpoints are sufficiently well-linked in their corresponding clusters. Chekuri and Chuzhoy [CC16] employed the following definition of well-linkedness:

Definition. We say that a set T of vertices is α -well-linked in a graph G , if for every partition (A, B) of the vertices of G into two subsets, $|E(A, B)| \geq \alpha \cdot \min\{|A \cap T|, |B \cap T|\}$.

The next observation follows immediately from Menger's theorem.

Observation 4.19 If a set T of vertices is edge-well-linked in a graph G , then T is 1-well-linked in G .

Next, we state the Boosting Theorems of [CC16].

Theorem 4.20 (Theorem 2.14 in [CC16]) Suppose we are given a connected graph $G = (V, E)$ with maximum vertex degree at most $\Delta \geq 3$ and a set $T \subseteq V$ of $\hat{\kappa}$ vertices, such that T is α -well-linked in G , for some $0 < \alpha \leq 1$. Then there is a subset $T' \subseteq T$ of $\lceil \frac{3\alpha\hat{\kappa}}{10\Delta} \rceil$ vertices, such that T' is node-well-linked in G .

Theorem 4.21 (Theorem 2.9 in [CC16]) Suppose we are given a graph G with maximum vertex degree at most Δ , and two disjoint subsets T_1, T_2 of vertices of G , with $|T_1|, |T_2| \geq \hat{\kappa}$, such that $T_1 \cup T_2$ is α -well-linked in G , for some $0 < \alpha \leq 1$, and each one of the sets T_1, T_2 is node-well-linked in G . Let $T'_1 \subseteq T_1, T'_2 \subseteq T_2$, be any pair of subsets with $|T'_1| = |T'_2| \leq \frac{\alpha\hat{\kappa}}{2\Delta}$. Then (T'_1, T'_2) are linked in G .

Recall that we are given a weak Path-of-Sets System $\mathbb{P} = (\mathcal{C}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, A_1, B_\ell)$ of length $\ell = g^2$ and width $w = g^2$, with $\mathcal{C} = (C_1, \dots, C_\ell)$, in a minor H'' of H , whose maximum vertex degree is bounded by 4. Consider some index $1 \leq i < \ell$, and recall that $B_i \subseteq C_i$ and $A_{i+1} \subseteq C_{i+1}$ are the sets of endpoints of the paths of \mathcal{P}_i lying in C_i and C_{i+1} , respectively. Applying Theorem 4.20 to graph C_i with $T = B_i$, we obtain a subset $\tilde{B}_i \subseteq B_i$ of at least $w' = 3w/40$ vertices, such that \tilde{B}_i is node-well-linked in C_i . Let $\tilde{\mathcal{P}}_i \subseteq \mathcal{P}_i$ be the set of paths originating at the vertices of \tilde{B}_i , and let $\tilde{A}_{i+1} \subseteq A_{i+1}$ be the set of their endpoints, lying in C_{i+1} . We then apply Theorem 4.20 to graph C_{i+1} , with the set $T = \tilde{A}_{i+1}$ of vertices, to obtain a collection $\tilde{A}'_{i+1} \subseteq \tilde{A}_{i+1}$ of at least $w'' = 3w'/40 = \Omega(w)$ vertices, such that \tilde{A}'_{i+1} is node-well-linked in C_{i+1} . Let $\tilde{\mathcal{P}}'_i \subseteq \tilde{\mathcal{P}}_i$ be the set of paths terminating at the vertices of \tilde{A}'_{i+1} . Finally, we select an arbitrary subset $\mathcal{P}'_i \subseteq \tilde{\mathcal{P}}'_i$ of $\tilde{w} = \lfloor w''/8 \rfloor = \Omega(w)$ paths, and we let $B'_i \subseteq \tilde{B}_i$ and $A'_{i+1} \subseteq \tilde{A}'_{i+1}$ be the sets of vertices where the paths of \mathcal{P}'_i originate and terminate, respectively.

As our last step, we apply Theorem 4.20 to graph C_1 and a set $T = A_1$ of vertices, to obtain a subset $\tilde{A}_1 \subseteq A_1$ of w' vertices that are node-well-linked in C_1 , and we select an arbitrary subset $A'_1 \subseteq \tilde{A}_1$ of \tilde{w} vertices of A_1 . We select a subset $B'_\ell \subseteq B_\ell$ of \tilde{w} vertices similarly.

For each $1 \leq i \leq \ell$, we are now guaranteed that each of the sets A'_i, B'_i is node-well-linked in C_i , and, from Theorem 4.21, (A'_i, B'_i) are linked in C_i . The final strong Path-of-Sets System is: $\mathbb{P}' = (\mathcal{C}, \{\mathcal{P}'_i\}_{i=1}^{\ell-1}, A'_1, B'_\ell)$; its length is $\ell = g^2$, and its width is $\tilde{w} = \Omega(g^2)$.

4.7 Parameters

For convenience, this subsection summarizes the main parameters used in Section 4.

- The cardinalities of the original sets A, B and X of vertices are $\kappa \geq 2^{22}g^{10} \log g$.

- The depth of the pseud-grid is $D = 64g^4$. We obtain D sets $\mathcal{R}_1, \dots, \mathcal{R}_D$ of paths, of cardinality at most g^2 each.
- The total number of paths in the set \mathcal{R} is N , where $D \leq N \leq g^2 D$, so $64g^4 \leq N \leq 64g^6$.
- The number of slices is $M = 8g^4 \log g$.
- For each slice i , we define subsets $\Sigma'_i \subseteq \Sigma_i$ and $\mathcal{Q}'_i \subseteq \mathcal{Q}_i$ of paths, such that $(\Sigma'_i, \mathcal{Q}'_i)$ are $(4g^2, D/2)$ -intersecting.
- After the well-linked decomposition, we obtain $M = 8g^4 \log g$ clusters - one cluster per slice. For each resulting cluster C_i , $|\tilde{\Sigma}_i| \geq D/4$, and the endpoints of the paths in $\tilde{\Sigma}_i$ are g^2 -weakly-well-linked in C_i .

5 Proof of Theorem 3.1

In this section we complete the proof of Theorem 3.1. Recall that we are given a graph H , and three disjoint subsets A, B, X of its vertices, each of cardinality $\kappa \geq 2^{22}g^9 \log g$, such that every vertex in X has degree 1. We are also given a set $\tilde{\mathcal{P}}$ of κ node-disjoint paths connecting vertices of A to vertices of B , and a set $\tilde{\mathcal{Q}}$ of κ node-disjoint paths connecting vertices of A to vertices of X . Our goal is to prove that either H contains an (A, B, X) -crossbar of width g^2 , or that its minor contains a strong Path-of-Sets System whose length and width are both at least $\Omega(g^2)$. We follow the framework of the proof in Section 4, but we introduce some changes that will allow us to save a factor of g on the parameter κ . The first three steps of the proof are identical to those in Section 4, except that we use slightly weaker parameters for the slicing.

Step 1: Pseudo-Grid

We start by defining the sets \mathcal{P}, \mathcal{Q} of κ node-disjoint paths exactly like in Section 4. We set the parameter $D = 64g^4$ like in Section 4. Note that the inequality $D \leq \kappa/(2g^2)$ still holds for the new value of κ , and that Theorem 4.1 holds regardless of the specific value of κ . We apply Theorem 4.1 to the graph H . If the outcome is an (A, B, X) -crossbar of width g^2 , then we return this crossbar and terminate the algorithm. Therefore, we assume from now on that the outcome of the theorem is a pseudo-grid of depth D . As before, we denote $N = |\mathcal{R}|$, so $D \leq N \leq g^2 D$ holds.

Step 2: Slicing

We define the subset $\mathcal{Q}'' \subseteq \mathcal{Q}'$ of paths exactly as before: Recall that for each $1 \leq i \leq D$, there are at most $2g^2$ paths $Q \in \mathcal{Q}'$, such that Q does not intersect any path of \mathcal{R}_i . We discard all such paths from \mathcal{Q}' , obtaining a set $\mathcal{Q}'' \subseteq \mathcal{Q}'$ of paths. Therefore, we discard at most $2g^2 D = 128g^6 < \kappa/8$ paths, and, since $\mathcal{Q}' = \lceil \kappa/4 \rceil$, we get that $|\mathcal{Q}''| \geq \kappa/8$. We are now guaranteed that Property I1 holds (that is, for each $1 \leq i \leq D$, each path $Q \in \mathcal{Q}''$ intersects some path of \mathcal{R}_i). As before, we denote by $A' \subseteq A$ and $B' \subseteq B$ the sets of endpoints of the paths of \mathcal{R} lying in A and B , respectively, and we let H' be the sub-graph of H , obtained by taking the union of all paths in \mathcal{R} and all paths in \mathcal{Q}'' .

We follow the algorithm from Section 4.2 to modify the graph H' and the sets $\mathcal{R}, \mathcal{Q}''$ of paths, such that the resulting graph H'' is a minor of H ; property I1 continues to hold; and H'' has the perfect unique linkage property with respect to A' and B' , with the unique linkage being \mathcal{R} .

Lastly, we perform a slicing of the paths in \mathcal{R} , by applying Theorem 4.5 to graph H'' , but the number of slices that we obtain is somewhat smaller. Specifically, we set $\hat{M} = 128g^3 \log g$ and $\hat{w} = 2^{11}g^6$. Notice that, since $N \leq g^2D = 64g^6$ and $|\mathcal{Q}''| \geq \kappa/8 \geq 2^{19}g^9 \log g$, $|\mathcal{Q}''| \geq \hat{M}\hat{w} + (\hat{M} + 1)\hat{N}$ holds. Therefore, we obtain an M_1 -slicing $\Lambda = \{\Lambda(R)\}_{R \in \mathcal{R}}$ of \mathcal{R} , of width w with respect to \mathcal{Q}'' , where $M_1 = 128g^3 \log g$, and $w = 2^{11}g^6$. As before, for every $1 \leq i \leq M_1$, we denote by $\Sigma_i = \{\sigma_i(R) \mid R \in \mathcal{R}\}$. We denote the subset $\hat{\mathcal{Q}}_i \subseteq \mathcal{Q}''$ of paths corresponding to Σ_i by \mathcal{Q}_i . We call $(\Sigma_i, \mathcal{Q}_i)$ the *i*th slice of Λ . Notice that so far we have followed the proof from Section 4 exactly, except that the number of slices that we obtain is smaller by factor g .

Step 3: Intersecting Path Sets

This step is also virtually the same as in Section 4. Let $\hat{w} = 4g^2$ and $\hat{D} = D/2$. Consider some $1 \leq i \leq M_1$, and the corresponding sets Σ_i, \mathcal{Q}_i of paths. Denote $\hat{\mathcal{R}} = \Sigma_i$ and $\hat{\mathcal{Q}} = \mathcal{Q}_i$. Recall that each path $Q \in \mathcal{Q}_i$ intersects at least D paths of Σ_i . It is also easy to verify that $|\hat{\mathcal{Q}}| \geq 2|\hat{\mathcal{R}}|\hat{w}\hat{D}$. Indeed, $|\hat{\mathcal{Q}}| \geq 2^{11}g^6$, while $2|\hat{\mathcal{R}}|\hat{w}\hat{D} = 16N \cdot g^2/D \leq 16g^4$. We can now apply Lemma 4.8 to sets $(\Sigma_i, \mathcal{Q}_i)$ of paths, with parameters \hat{D}, \hat{w} , to obtain a partition (Σ'_i, Σ''_i) of Σ_i , and a subset $\mathcal{Q}'_i \subseteq \mathcal{Q}_i$ of paths, such that $(\Sigma'_i, \mathcal{Q}'_i)$ is a $(4g^2, D/2)$ -intersecting pair of path sets, $|\mathcal{Q}'_i| \geq 2^{10}g^6$, and each path of Σ''_i intersects at most $4g^2$ paths of \mathcal{Q}'_i .

So far our proof followed exactly the proof from Section 4. Unfortunately, the number of slices M_1 that we obtain is smaller than the one in Section 4 by a factor of g . In order to get around this problem, we distinguish between two cases. Intuitively, in one of the cases (Case 2) we will be able to increase the number of slices so that we can employ the rest of the proof from Section 4. In the other case (Case 1) the number of slices will stay low, but we will compensate this by ensuring that the cardinalities of the sets Σ'_i are large in many of the slices.

For $1 \leq i \leq M_1$, we say that slice $(\Sigma_i, \mathcal{Q}_i)$ is of type 1 iff $|\Sigma'_i| \geq N/g$, and we say that it is of type 2 otherwise. We say that Case 1 happens if at least half the slices are of type 1; otherwise, we say that Case 2 happens.

5.1 Case 1: Many Type-1 Slices

If Case 1 happens, at least half the slices are of type 1. We will ignore all type-2 slices, so we can assume that in Case 1 the number of slices is $64g^3 \log g$, that we denote, abusing the notation, by M_1 , and that in every slice $(\Sigma_i, \mathcal{Q}_i)$, $|\Sigma'_i| \geq N/g$. Recall that for each i , $(\Sigma'_i, \mathcal{Q}'_i)$ are $(4g^2, D/2)$ -intersecting, and $|\mathcal{Q}'_i| \geq 2^{10}g^6$.

Step 4: Well-Linked Decomposition

Fix some index i , such that $(\Sigma_i, \mathcal{Q}_i)$ is a type-1 slice. As before, we let H'_i be a graph obtained from the union of the paths in Σ_i and \mathcal{Q}_i . And we denote $\hat{D} = D/4$ and $\hat{w} = g^2$, so that $(\Sigma'_i, \mathcal{Q}'_i)$ are $(4\hat{w}, 2\hat{D})$ -intersecting, and $\hat{D} \geq 8\hat{w}$, since $D = 64g^4$. As in Section 4.4, we apply Theorem 4.11 to graph H'_i , with $\hat{\Sigma} = \Sigma'_i$, $\hat{\mathcal{Q}} = \mathcal{Q}'_i$, and parameters \hat{D}, \hat{w} . Let \mathcal{C}_i be the resulting collection of happy clusters. For each such cluster $C \in \mathcal{C}_i$, we denote $\tilde{\Sigma}(C) = \hat{\Sigma}(C)$. As before, the endpoints of the paths of $\tilde{\Sigma}(C)$ are g^2 -weakly well-linked in C , and $|\tilde{\Sigma}(C)| \geq D/4$. We denote $\tilde{\Sigma}_i = \bigcup_{C \in \mathcal{C}_i} \tilde{\Sigma}(C)$. Notice that Theorem 4.11 guarantees that $|\tilde{\Sigma}_i| \geq |\Sigma'_i|/4 \geq N/(4g)$. Let \mathcal{C} be the union of all sets \mathcal{C}_i , where i is a type-1 slice, and let $\tilde{\Sigma}$ be the union of all corresponding sets $\tilde{\Sigma}_i$, so $|\tilde{\Sigma}| \geq M_1N/4g$.

Intuitively, for some of the clusters $C \in \mathcal{C}$, $|\tilde{\Sigma}(C)|$ may be very large, and then it is sufficient for us to

have a small number of such clusters. But it is possible that for most clusters in C , $|\tilde{\Sigma}(C)|$ is relatively small – possibly as small as $D/4$. In this case, it would be helpful for us to argue that the number of such clusters is large. In other words, we would like to obtain a tradeoff between the number of clusters C and the cardinality of their corresponding path sets $|\tilde{\Sigma}(C)|$.

In order to do so, we group the clusters geometrically. Recall that for each $C \in \mathcal{C}$, $D/4 \leq |\tilde{\Sigma}(C)| \leq N \leq g^2 D$. For $0 \leq j < 2 \log g + 2$, we say that cluster C belongs to class \mathcal{S}_j iff $D \cdot 2^j / 4 \leq |\tilde{\Sigma}(C)| < D \cdot 2^{j+1} / 4$. If $C \in \mathcal{S}_j$, then we say that all paths in $\tilde{\Sigma}(C)$ belong to class j . Then there must be an index j , such that the number of paths in $\tilde{\Sigma}$ that belong to class j is at least $\frac{|\tilde{\Sigma}|}{2 \log g + 2} \geq \frac{M_1 N}{16g \log g}$. We let $\tilde{\mathcal{C}}' = \mathcal{S}_j$, and we let $\tilde{\Sigma}' \subseteq \tilde{\Sigma}$ be the set of all paths that belong to class j . Note that $|\tilde{\mathcal{C}}'| \geq \frac{|\tilde{\Sigma}'|}{D \cdot 2^{j+1} / 4} \geq \frac{M_1 N}{4D \cdot 2^{j+1} g \log g} = \frac{64g^3 N \log g}{8D \cdot 2^j g \log g} = \frac{8Ng^2}{D \cdot 2^j}$, while for each $C \in \mathcal{C}$, $|\tilde{\Sigma}(C)| \geq D \cdot 2^j / 4$.

To summarize the steps in Case 1 until now, we have obtained a collection of at least $\frac{8Ng^2}{D \cdot 2^j}$ clusters, that we denote, abusing the notation, by $\tilde{\mathcal{C}}$. For each slice i , we may have a number of clusters of $\tilde{\mathcal{C}}$ that belong to that slice. We denote the set of all such clusters by $\tilde{\mathcal{C}}_i$. Notice however that, if $C, C' \in \tilde{\mathcal{C}}_i$ and $C \neq C'$, then $\tilde{\Sigma}(C) \cap \tilde{\Sigma}(C') = \emptyset$. Moreover, if $\mathcal{R}' \subseteq \mathcal{R}$ is the set of all paths containing the segments of $\tilde{\Sigma}(C)$, and \mathcal{R}'' is defined similarly for C' , then $\mathcal{R}' \cap \mathcal{R}'' = \emptyset$. We are guaranteed that for each cluster $C \in \tilde{\mathcal{C}}$, $|\tilde{\Sigma}(C)| \geq D \cdot 2^j / 4$, and the endpoints of the paths of $\tilde{\Sigma}(C)$ are g^2 -weakly well-linked in C .

Step 5: a Path-of-Sets System

In this step, we employ Theorem 4.15 to construct a weak Path-of-Sets System of length and width g^2 in H'' . We use the parameters $\hat{M} = |\tilde{\mathcal{C}}| \geq \frac{8Ng^2}{D \cdot 2^j}$, $\hat{w} = g^2$, $\hat{N} = N$, and $\hat{D} = \frac{D \cdot 2^j}{4}$. Recall that set \mathcal{C} may contain clusters from the same slice. We order the clusters in \mathcal{C} as follows: first, we order the clusters in the increasing order of their slices; the clusters inside the same slice are ordered arbitrarily. Let $C_1, \dots, C_{\hat{M}}$ be the resulting ordering of the clusters. We define the sets $S_1, \dots, S_{\hat{M}}$ exactly as before. An important observation is that, if C_i, C_j belong to the same slice, then $S_i \cap S_j = \emptyset$.

We need to verify that the conditions of Theorem 4.15 hold for our choice of parameters. The first condition is that $\hat{N} \geq 3\hat{w}$. Since we use the same parameters \hat{N} and \hat{w} as in Section 4.5, this inequality continues to hold. The second condition is $\hat{D}^2 \geq 4\hat{N}\hat{w}$. Recall that this condition held for the parameters in Section 4, where the values of \hat{N} and \hat{w} were the same, and $\hat{D} = D/4$ was smaller than the current value of \hat{D} , so this condition continues to hold. The third condition is that $\hat{D}\hat{M} \geq 2\hat{N}\hat{w}$. This condition is easy to verify by substituting the values of the relevant parameters.

We now use Theorem 4.15 to obtain a sequence $1 \leq i_1 < i_2 < \dots < i_{g^2} \leq \hat{M}$ of indices, such that for all $1 \leq j < g^2$, $|S_{i_j} \cap S_{i_{j+1}}| \geq g^2$. Notice that each resulting cluster $C_{i_1}, C_{i_2}, \dots, C_{i_{g^2}}$ must belong to a different slice. The remainder of the construction of the weak Path-of-Sets system is done exactly as in subsection 4.5. Therefore, we obtain a weak Path-of-Sets system of width g^2 and length g^2 . Our last step is to convert it into a strong Path-of-Sets system using the same procedure as in Section 4.6.

5.2 Case 2: Many Type-2 Slices

Assume now that Case 2 happens, so we have at least $\frac{1}{2}M_1 = 64g^3 \log g$ type-2 slices $\{(\Sigma_i, \mathcal{Q}_i)\}_i$. Recall that for each i , we are given a partition (Σ'_i, Σ''_i) of Σ_i , and a subset $\mathcal{Q}'_i \subseteq \mathcal{Q}_i$ of at least $2^{10}g^6$ paths, such that $(\Sigma'_i, \mathcal{Q}'_i)$ are $(4g^2, D/2)$ -intersecting, and every path in Σ''_i intersects at most $4g^2$ paths of \mathcal{Q}'_i . Recall that in each type-2 slice $(\Sigma_i, \mathcal{Q}_i)$, $|\Sigma'_i| < N/g$.

Intuitively, we are interested in either obtaining a large number of slices, or in obtaining a large number

of paths in the sets Σ'_i of each such slice. Type-1 slices achieve the latter. But in type-2 slices, the cardinalities of sets Σ'_i are too small for us. Fortunately, we can exploit this fact in order to increase the number of slices.

Theorem 5.1 *Assume that Case 2 happens, and that H does not contain an (A, B, X) -crossbar of width g^2 . Then there is an M_2 -slicing of \mathcal{R} , of width at least N/g with respect to \mathcal{Q}'' , where $M_2 = 8g^4 \log g$.*

Proof: The proof directly follows from the following lemma.

Lemma 5.2 *Assume that H does not contain an (A, B, X) -crossbar of width g^2 . Let $1 \leq i \leq M_1$ be an index, such that $(\Sigma_i, \mathcal{Q}_i)$ is a type-2 slice of the original slicing Λ . Then there is an \hat{M} -slicing Λ_i of Σ_i of width at least N/g with respect to \mathcal{Q}_i , where $\hat{M} = g$.*

Before we provide the proof of Lemma 5.2, we complete the proof of Theorem 5.1 using it. Recall that in Case 2, there are at least $64g^3 \log g$ slices of Λ of type-2 slices. Each such slice $(\Sigma_i, \mathcal{Q}_i)$ is then in turn sliced into g slices using Lemma 5.2. Therefore, by combining the slicing Λ together with the individual slicings Λ_i for all type-2 slices $(\Sigma_i, \mathcal{Q}_i)$, we obtain a slicing of \mathcal{R} , where the number of slices is at least $(64g^3 \log g) \cdot g > 8g^4 \log g = M_2$. The width of the new slicing with respect to \mathcal{Q}'' is at least N/g .

From now on we focus on the proof of Lemma 5.2. Let $1 \leq i \leq M_1$ be an index, such that $(\Sigma_i, \mathcal{Q}_i)$ is a type-2 slice of the original slicing Λ . Our goal is to produce an \hat{M} -slicing Λ_i of Σ_i of width at least N/g with respect to \mathcal{Q}_i , where $\hat{M} = g$. The idea is that we will discard all paths in $\mathcal{Q}_i \setminus \mathcal{Q}'_i$; ignore the paths in Σ''_i (for each such path $\sigma \in \Sigma''_i$ we will eventually produce a trivial slicing, where $v_0(\sigma)$ is the first endpoint of σ , and $v_1(\sigma) = v_2(\sigma) = \dots = v_{\hat{M}}(\sigma)$ is its last endpoint), and will focus on slicing the paths of Σ'_i , trying to achieve a slicing whose width is at least N/g with respect to \mathcal{Q}'_i . Unfortunately, some of the paths in \mathcal{Q}'_i may intersect the paths of Σ''_i , so our first step is to get rid of all such intersections. We do so using the following claim. Recall that $|\mathcal{Q}'_i| \geq 2^{10}g^6$.

Claim 5.3 *Assume that H does not contain an (A, B, X) -crossbar of width g^2 , and let $(\Sigma_i, \mathcal{Q}_i)$ be a slice of type 2. Then at least 2^9g^6 paths in the set \mathcal{Q}'_i are disjoint from the paths in Σ''_i .*

Proof: Assume otherwise. Let $\mathcal{B} \subseteq \mathcal{Q}'_i$ be the set of all paths that have non-empty intersection with paths in Σ''_i , so $|\mathcal{B}| \geq 2^9g^6$. We further partition the set \mathcal{B} into two subsets: set \mathcal{B}_1 contains all paths $Q \in \mathcal{B}$ that intersect at least $8g^2$ paths of Σ''_i , and $\mathcal{B}_2 = \mathcal{B} \setminus \mathcal{B}_1$.

We first show that $|\mathcal{B}_1| \leq 32g^6$. Recall that $|\Sigma''_i| \leq N \leq 64g^6$, and each path $\sigma \in \Sigma''_i$ intersects at most $4g^2$ paths of \mathcal{Q}'_i . Therefore, there are at most $256g^8$ pairs (σ, Q) of paths, with $\sigma \in \Sigma''_i$ and $Q \in \mathcal{B}$, such that $\sigma \cap Q \neq \emptyset$. As each path of \mathcal{B}_1 intersects at least $8g^2$ paths of Σ''_i , we get that $|\mathcal{B}_1| \leq 256g^8/8g^2 \leq 32g^6$.

We conclude that $|\mathcal{B}_2| \geq 2^8g^6$. We exploit this fact to construct a (A, B, X) -crossbar of width g^2 in H . In order to do so, we perform g^2 iterations, where in each iteration we add some path P connecting a vertex of A to a vertex of B to the crossbar, and its corresponding path Q_P , thus iteratively constructing the crossbar $(\mathcal{P}^*, \mathcal{Q}^*)$. In every iteration, we will delete some paths from Σ''_i and from \mathcal{B}_2 . The path P that we add to \mathcal{P}^* is a path from \mathcal{R} , that contains some segment of Σ''_i , and its corresponding path Q_P is a sub-path of a path in \mathcal{B}_2 . We start with $\mathcal{P}^*, \mathcal{Q}^* = \emptyset$, and we maintain the following invariants:

- All paths in the current sets $\mathcal{P}^*, \mathcal{Q}^*$ are disjoint from all paths in $\Sigma''_i, \mathcal{B}_2$; moreover, each path $R \in \mathcal{R}$ that contains a segment $\sigma \in \Sigma''_i$ is disjoint from all paths in $\mathcal{P}^* \cup \mathcal{Q}^*$; and

- Each remaining path in \mathcal{B}_2 intersects some path in the remaining set Σ_i'' .

At the beginning, $\mathcal{P}^*, \mathcal{Q}^* = \emptyset$, and the invariants hold. Assume now that the invariants hold at the beginning of iteration j . The j th iteration is executed as follows. We let $Q \in \mathcal{B}_2$ be any path, and we let $\sigma \in \Sigma_i''$ be any path intersecting Q . We add the unique path $P \in \mathcal{R}$ that contains σ to \mathcal{P}^* , and we add a sub-path of Q , connecting a vertex of P to a vertex of X to \mathcal{Q}^* , as Q_P (recall that by the definition of a pseudo-grid, every path in \mathcal{Q}' has an endpoint that lies in X)⁴. Let \mathcal{S}_j be the collection of all paths of the current set Σ_i'' that intersect Q , so $|\mathcal{S}_j| < 8g^2$. Let $\mathcal{Y}_j \subseteq \mathcal{B}_2$ be the set of all paths that intersect the paths of \mathcal{S}_j , so $|\mathcal{Y}_j| \leq |\mathcal{S}_j| \cdot 4g^2 \leq 32g^4$. We delete the paths of \mathcal{S}_j from Σ_i'' , and we delete from \mathcal{B}_2 all paths of \mathcal{Y}_j . It is easy to verify that the invariants continue to hold after this iteration (for the second invariant, recall that each path of the original set \mathcal{B}_2 intersected some path of Σ_i'' ; whenever a path σ is deleted from Σ_i'' , we delete all paths that intersect it from \mathcal{B}_2 . Therefore, each path that remains in \mathcal{B}_2 must intersect some path that remains in Σ_i''). In every iteration, at most $32g^4$ paths are deleted from \mathcal{B}_2 , while at the beginning $|\mathcal{B}_2| \geq 2^8 g^6$. Therefore, we can carry this process for g^2 iterations, after which we obtain an (A, B, X) -crossbar of width g^2 . \square

We let $\tilde{\mathcal{Q}}_i \subseteq \mathcal{Q}'_i$ be the set of at least $512g^6$ paths of \mathcal{Q}'_i that are disjoint from the paths of Σ_i'' .

Let A_i be the set of vertices that serve as the first endpoint of the paths in Σ_i , and let $A'_i \subseteq A_i$ be defined similarly for Σ'_i . Similarly, let B_i be the set of vertices that serve as the last endpoint of the paths in Σ_i , and let $B'_i \subseteq B_i$ be defined similarly for Σ'_i . We let H_i be the graph obtained from the union of the paths in Σ_i and \mathcal{Q}_i , and we let H'_i be the graph obtained from the union of the paths in Σ'_i and $\tilde{\mathcal{Q}}_i$.

Observation 5.4 *Graph H'_i has the perfect unique linkage property with respect to (A'_i, B'_i) , with the unique linkage being Σ'_i .*

Assume first that the observation is true. Recall that $|\Sigma'_i| \leq N/g$, and $|\tilde{\mathcal{Q}}_i| \geq 512g^6$. We now invoke Theorem 4.5 with $\hat{\mathcal{R}} = \Sigma'_i$, $\hat{\mathcal{Q}} = \tilde{\mathcal{Q}}_i$, $\hat{N} = |\Sigma'_i| \leq N/g$; $\hat{M} = g$ and $\hat{w} = N/g$, to obtain an \hat{M} -slicing of Σ'_i of width N/g with respect to $\tilde{\mathcal{Q}}_i$. In order to do so, we need to verify that $|\tilde{\mathcal{Q}}_i| \geq \hat{M}\hat{w} + (\hat{M} + 1)|\Sigma'_i|$. But $\hat{M}\hat{w} + (\hat{M} + 1)|\Sigma'_i| \leq (2\hat{M} + 1)N/g \leq 6N$, while $|\tilde{\mathcal{Q}}_i| \geq 512g^6 \geq 6N$, as $N \leq Dg^2 = 64g^6$. We conclude that there exists an \hat{M} slicing Λ_i of Σ'_i , whose width with respect to $\tilde{\mathcal{Q}}_i$ is N/g . We extend this slicing to an \hat{M} -slicing of Σ_i in a trivial way: for every path $\sigma \in \Sigma_i''$, we let $v_0(\sigma)$ be its first endpoint, and we let $v_1(\sigma) = v_2(\sigma) = \dots = v_{\hat{M}}(\sigma)$ be its last endpoint. Since the paths of $\tilde{\mathcal{Q}}_i$ are disjoint from the paths of Σ_i'' , this defines an \hat{M} -slicing of Σ_i of width N/g with respect to $\tilde{\mathcal{Q}}_i$, and hence with respect to \mathcal{Q}_i as well. It now remains to prove Observation 5.4.

Proof of Observation 5.4. We first claim that H_i has the unique perfect linkage property with respect to (A_i, B_i) , with the unique linkage being Σ_i . Indeed, for each path $R \in \mathcal{R}$, no vertex of $R \setminus \sigma_i(R)$ belongs to H_i . If there is a different (A_i, B_i) linkage $\tilde{\Sigma}_i \neq \Sigma_i$, then we could replace Σ_i with $\tilde{\Sigma}_i$ in \mathcal{R} , obtaining a different (A', B') -linkage in H'' , violating Observation 4.4. It is also immediate to verify that every vertex of H_i should lie on some path in Σ_i , as \mathcal{R} has the perfect unique linkage property in H'' .

Consider now the graph H_i , after we delete all paths of $\mathcal{Q}_i \setminus \tilde{\mathcal{Q}}_i$ from it. In this new graph, no path of Σ'_i may lie in the same connected component as a path of Σ_i'' , as we have deleted all paths that may intersect the paths of Σ_i'' . Clearly, this new graph still has the perfect unique linkage property with respect to (A_i, B_i) , with the unique linkage being Σ_i . Graph H'_i is obtained from H_i by deleting all

⁴In fact we have contracted edges on the paths in \mathcal{Q}' when constructing the graph H'' in Step 2, in order to ensure the perfect unique linkage property; formally, in order to obtain the path Q_P , we need to un-contract the path Q and then take a segment of the resulting path connecting a vertex of P to a vertex of X .

paths of Σ_i'' from it. Since each of these paths lies in a distinct connected component, it is easy to verify that H_i' must have the perfect unique linkage property with respect to (A_i', B_i') , with the unique linkage being Σ_i' . \square \square

If Theorem 5.1 returns an (A, B, X) -crossbar of width g^2 , then we output this crossbar and terminate the algorithm. Therefore, we assume from now on that the theorem returns an M_2 -slicing of \mathcal{R} , whose width with respect to \mathcal{Q}'' is at least N/g . We will ignore the original slicing for Case 2, and will denote this new slicing by Λ . Abusing the notation, we denote, for each $1 \leq i \leq M_2$, the i th slice of this new slicing by $(\Sigma_i, \mathcal{Q}_i)$, where $\mathcal{Q}_i \subseteq \mathcal{Q}''$ and Σ_i contains the i th segment of each path $R \in \mathcal{R}$.

As before, for each $1 \leq i \leq M_2$, we employ Lemma 4.8 in order to partition the set Σ_i into two subsets, Σ_i', Σ_i'' , and compute a subset $\mathcal{Q}_i' \subseteq \mathcal{Q}_i$ of paths, such that $(\Sigma_i', \mathcal{Q}_i')$ are $(4g^2, D/2)$ -intersecting path sets and $|\mathcal{Q}_i'| \geq N/(2g)$. In order to do so, we set $\hat{w} = 4g^2$ and $\hat{D} = D/2$, and we denote $\hat{\mathcal{R}} = \Sigma_i$ and $\hat{\mathcal{Q}} = \mathcal{Q}_i$. As before, every path of \mathcal{Q}_i intersects at least $D = 2\hat{D}$ paths of Σ_i .

We now to verify that $|\hat{\mathcal{Q}}| \geq \frac{2|\hat{\mathcal{R}}|\hat{w}}{\hat{D}}$. Indeed, $|\hat{\mathcal{Q}}| = |\mathcal{Q}_i| \geq 2|\Sigma_i| \cdot 4g^2/(D/2) = 16Ng^2/D$, as $D = 64g^4$. We can now use Lemma 4.8 to obtain a partition (Σ_i', Σ_i'') of Σ_i into two subsets, and a subset $\mathcal{Q}_i' \subseteq \mathcal{Q}_i$ of paths, such that $(\Sigma_i', \mathcal{Q}_i')$ are $(4g^2, D/2)$ -intersecting path sets and $|\mathcal{Q}_i'| \geq N/(2g)$.

The remainder of the proof repeats the Steps 4–6 from the proof in Section 3. Observe that our starting point is now identical to the starting point of Step 4 in Section 4. The values of parameters D and N remained unchanged. We have computed a slicing Λ of \mathcal{R} into $M_2 = 8g^4 \log g$ slices - which is equal to the number of slices used in Section 4. We have also computed, for each $1 \leq i \leq M_2$, subsets $\Sigma_i' \subseteq \Sigma_i$, and $\mathcal{Q}_i' \subseteq \mathcal{Q}_i$ of node-disjoint paths, such that $(\Sigma_i', \mathcal{Q}_i')$ are $(4g^2, D/2)$ -intersecting. Therefore, we can now repeat Steps 4–6 from the proof in Section 3 to obtain a strong Path-of-Sets system of length $\Omega(g^2)$ and width $\Omega(g^2)$ in H'' .

5.3 Parameters

For convenience, this subsection summarizes the main parameters used in Section 5.

- The cardinalities of the original sets A, B and X of vertices from the statement of Theorem 3.1 are $\kappa \geq 2^{22}g^9 \log g$.
- The depth of the pseudo-grid is $D = 64g^4$. We obtain D sets $\mathcal{R}_1, \dots, \mathcal{R}_D$ of paths, of cardinality at most g^2 each.
- The total number of paths in the set \mathcal{R} is N , where $D \leq N \leq g^2D$, so $64g^4 \leq N \leq 64g^6$.
- The number of slices in the first slicing is $M_1 = 128g^3 \log g$. The width of the slicing is $w = 2^{11}g^6$. For each type-1 slice i , $|\Sigma_i'| \geq N/g$.
- The number of slices in the second slicing is $M_2 = 8g^4 \log g$, and the width is $w_2 = N/g$.
- In both Case 1 and Case 2, for each slice i , we define subsets $\Sigma_i' \subseteq \Sigma_i$ and $\mathcal{Q}_i' \subseteq \mathcal{Q}_i$ of paths, such that $(\Sigma_i', \mathcal{Q}_i')$ are $(4g^2, D/2)$ -intersecting.

6 Acknowledgements

Part of the work was done while the first author was a Weston visiting professor in the Department of Computer Science and Applied Mathematics, Weizmann Institute. The first author was supported

in part by NSF grants CCF-1318242 and CCF-1616584. The second author was supported in part by NSF grants CCF-1318242 and CCF-1616584.

References

- [CC15] Chandra Chekuri and Julia Chuzhoy. Degree-3 treewidth sparsifiers. In Piotr Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 242–255. SIAM, 2015.
- [CC16] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM (JACM)*, 63(5):40, 2016.
- [Chu15] Julia Chuzhoy. Excluded grid theorem: Improved and simplified. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC '15*, pages 645–654, New York, NY, USA, 2015. ACM.
- [Chu16] Julia Chuzhoy. Improved bounds for the excluded grid theorem. *arXiv preprint arXiv:1602.02629*, 2016.
- [CK14] Chandra Chekuri and Nitish Korula. A graph reduction step preserving element-connectivity and packing steiner trees and forests. *SIAM J. Discrete Math.*, 28(2):577–597, 2014.
- [CN18] Julia Chuzhoy and Rachit Nimavat. Large minors in expanders. Manuscript, 2018.
- [DH07a] E.D. Demaine and M Hajiaghayi. The Bidimensionality Theory and Its Algorithmic Applications. *The Computer Journal*, 51(3):292–302, November 2007.
- [DH07b] Erik D Demaine and MohammadTaghi Hajiaghayi. Quickly deciding minor-closed parameters in general graphs. *European Journal of Combinatorics*, 28(1):311–314, January 2007.
- [DHK09] Erik D Demaine, MohammadTaghi Hajiaghayi, and Ken-ichi Kawarabayashi. Algorithmic graph minor theory: Improved grid minor bounds and Wagner’s contraction. *Algorithmica*, 54(2):142–180, 2009.
- [FST11] Fedor V Fomin, Saket Saurabh, and Dimitrios M Thilikos. Strengthening Erdős-Pósa property for minor-closed graph classes. *Journal of Graph Theory*, 66(3):235–240, 2011.
- [HO96] H. R. Hind and O. Oellermann. Menger-type results for three or more vertices. *Congressus Numerantium*, 113:179–204, 1996.
- [KK12] Ken-ichi Kawarabayashi and Yusuke Kobayashi. Linear min-max relation between the treewidth of h-minor-free graphs and its largest grid. In *STACS’12 (29th Symposium on Theoretical Aspects of Computer Science)*, volume 14, pages 278–289. LIPIcs, 2012.
- [KKOV07] Rohit Khandekar, Subhash Khot, Lorenzo Orecchia, and Nisheeth K Vishnoi. On a cut-matching game for the sparsest cut problem. *Univ. California, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2007-177*, 2007.
- [KR96] Jon Kleinberg and Ronitt Rubinfeld. Short paths in expander graphs. In *In Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 86–95, 1996.

- [Kri18] Michael Krivelevich. Expanders - how to find them, and what to find in them. *arXiv preprint arXiv:1812.11562*, 2018.
- [KRV09] Rohit Khandekar, Satish Rao, and Umesh Vazirani. Graph partitioning using single commodity flows. *J. ACM*, 56(4):19:1–19:15, July 2009.
- [LS15] Alexander Leaf and Paul Seymour. Tree-width and planar minors. *Journal of Combinatorial Theory, Series B*, 111:38–53, 2015.
- [Mad78] W. Mader. A reduction method for edge connectivity in graphs. *Ann. Discrete Math.*, 3:145–164, 1978.
- [OSVV08] L. Orecchia, L.J. Schulman, U.V. Vazirani, and N.K. Vishnoi. On partitioning graphs via single commodity flows. In *Proc. of ACM STOC*, pages 461–470. ACM, 2008.
- [Ree97] Bruce Reed. *Surveys in Combinatorics*, chapter Treewidth and Tangles: A New Connectivity Measure and Some Applications. London Mathematical Society Lecture Note Series. Cambridge University Press, 1997.
- [RS83] Neil Robertson and Paul D Seymour. Graph minors. I. Excluding a forest. *Journal of Combinatorial Theory, Series B*, 35(1):39–61, 1983.
- [RS86] Neil Robertson and P D Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, August 1986.
- [RS95] Neil Robertson and Paul D Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of combinatorial theory, Series B*, 63(1):65–110, 1995.
- [RST94] N Robertson, P Seymour, and R Thomas. Quickly Excluding a Planar Graph. *Journal of Combinatorial Theory, Series B*, 62(2):323–348, November 1994.
- [SL15] Mohit Singh and Lap Chi Lau. Approximating minimum bounded degree spanning trees to within one of optimal. *J. ACM*, 62(1):1:1–1:19, March 2015.
- [Tho88] Carsten Thomassen. On the presence of disjoint subgraphs of a specified type. *Journal of Graph Theory*, 12(1):101–111, 1988.

A Proofs Omitted from Section 2

A.1 Proof of Claim 2.2

In order to prove the claim, we only need to show how to construct, for each $1 \leq i < \lceil \ell/2 \rceil$, a set $\hat{\mathcal{P}}_i$ of w' node-disjoint paths, connecting the vertices of B'_{2i-1} to the vertices of A'_{2i+1} , so that the paths of $\bigcup_{i=1}^{\lceil \ell/2 \rceil} \hat{\mathcal{P}}_i$ are disjoint from each other, and they are internally disjoint from $\bigcup_{j=1}^{\lceil \ell/2 \rceil} V(C'_j) = \bigcup_{j=1}^{\lceil \ell/2 \rceil} V(C_{2j-1})$. In order to do so, it is enough to show that for each $1 \leq i < \lceil \ell/2 \rceil$, there is a set $\hat{\mathcal{P}}_i$ of w' node-disjoint paths, connecting the vertices of B'_{2i-1} to the vertices of A'_{2i+1} , such that the paths of $\hat{\mathcal{P}}_i$ are contained in the graph $\mathcal{P}_{2i-1} \cup C_{2i} \cup \mathcal{P}_{2i}$.

Fix some $1 \leq i < \lceil \ell/2 \rceil$. Let $\hat{B} = B'_{2i-1}$, and let $\hat{A} = A'_{2i+1}$. Let $\tilde{\mathcal{P}}_1 \subseteq \mathcal{P}_{2i-1}$ be the set of paths originating from the vertices of \hat{B} , and let \hat{X} be the set of endpoints of the paths of $\tilde{\mathcal{P}}_1$ that lie in C_{2i} (see Figure A.1). Similarly, let $\tilde{\mathcal{P}}_2 \subseteq \mathcal{P}_{2i}$ be the set of paths terminating at the vertices of \hat{A} , and let \hat{Y} be the set of endpoints of these paths that lie in C_{2i} . Clearly, $\hat{X} \subseteq A_{2i}$ and $\hat{Y} \subseteq B_{2i}$. Since (A_{2i}, B_{2i})

are linked in C_{2i} , there is a set \mathcal{Q} of w' node-disjoint paths connecting \hat{X} to \hat{Y} in C_{2i} . By combining $\tilde{\mathcal{P}}_1, \mathcal{Q}, \tilde{\mathcal{P}}_2$, we obtain the desired set $\tilde{\mathcal{P}}_i$ of w' node-disjoint paths connecting \hat{B} to \hat{A} .

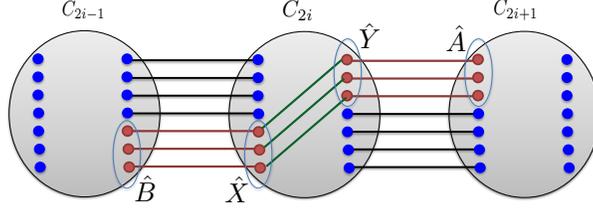


Figure 4: Stitching in a Path-of-Sets System. The paths of $\tilde{\mathcal{P}}_1$ and $\tilde{\mathcal{P}}_2$ are shown in red, the paths of \mathcal{Q} in green.

A.2 Proof of Theorem 2.3

Our starting point is the following two theorems, that were proved in [CC15] and [CC16], respectively.

Theorem A.1 (Theorem 1.1 in [CC15]) *Let G be a graph of treewidth k . Then there is a subgraph G' of G , whose maximum vertex degree is 3, and $\text{tw}(G') = \Omega(k/\text{poly log } k)$.*

Theorem A.2 (Theorem 3.4 in [CC16]) *There are constants $\hat{c}, \hat{c}' > 0$, such that for all integers $\ell, w, k > 1$ with $k/\log^{\hat{c}'} k > \hat{c}w\ell^{48}$, every graph G of treewidth at least k contains a strong Path-of-Sets System of length ℓ and width w .*

Let G be a graph of treewidth at least k . We use Theorem A.1 to obtain a subgraph $G' \subseteq G$ of treewidth $k' = \Omega(k/\text{poly log } k)$ and maximum vertex degree 3. Let $\ell' = 2\ell$ and let $w' = c^* \cdot w$, for a large enough constant c^* , that will be determined later. By appropriately setting the values of the constants c and c' in the statement of Theorem 2.3, we can ensure that $k'/\log^{\hat{c}'} k' > \hat{c}w'(\ell')^{48}$. From Theorem A.2, graph G' contains a strong Path-of-Sets System of length ℓ' and width w' . Our last step is to turn it into a hairy Path-of-Sets System of length ℓ and width w , using the following theorem, that was proved in [Chu16].

Theorem A.3 (Theorem 6.3 in [Chu16]) *For every integer $\Delta > 0$, there is an integer $c_\Delta > 0$ depending only on Δ , such that the following holds. Let G be any graph of maximum vertex degree at most Δ , and let A, B be two disjoint subsets of vertices of G , with $|A| = |B| = \kappa$, such that A and B are each node-well-linked in G , and (A, B) are node-linked in G . Then there are two disjoint clusters $C', S' \subseteq G$, a set \mathcal{Q} of at least κ/c_Δ node-disjoint paths connecting vertices of C' to vertices of S' , so that the paths of \mathcal{Q} are internally disjoint from $C' \cup S'$, and two subsets $A' \subseteq A \cap C', B' \subseteq B \cap C'$ of at least κ/c_Δ vertices each such that, if we denote by X' and Y' the endpoint of the paths of \mathcal{Q} lying in C' and S' respectively, then:*

- set Y' is node-well-linked in S' ;
- each of the three sets A', B' and X' is node-well-linked in C' ; and
- every pair of sets in $\{A', B', X'\}$ is node-linked in C' .

Using the above theorem, we show that any Path-of-Sets System can be transformed into a hairy Path-of-Sets System of roughly the same length and width, in the following lemma.

Lemma A.4 *Let G' be a graph of maximum vertex degree 3, and assume that for some $\ell', w' > 0$, G contains a strong Path-of-Sets System of length ℓ' and width w' . Then G contains a hairy Path-of-Sets System with length at least $\ell'/2$ and width at least $w'/(3c_\Delta)$, where c_Δ is the constant from Theorem A.3.*

Setting the constant c^* from the definition of w' to be $3c_\Delta$, from the above lemma, graph G' contains a hairy Path-of-Sets System of length at least $\ell = \ell'/2$ and width at least $w = w'/(3c_\Delta)$, completing the proof of Theorem 2.3. It now remains to prove Lemma A.4.

Proof of Lemma A.4. Let $\mathbb{P} = (\mathcal{S}, \{\mathcal{P}_i\}_{i=1}^{\ell'-1}, A_1, B_{\ell'})$ be the given strong Path-of-Sets System in G' , of length ℓ' and width w' . Recall that the maximum vertex degree in G' is 3. Let $1 \leq i \leq \ell'$ be an odd integer. We apply Theorem A.3 to graph C_i , with $A = A_i$ and $B = B_i$. We denote the resulting two clusters C' and S' by C'_i and S'_i , respectively, and we denote the resulting subsets A', B', X', Y' of vertices by A''_i, B''_i, X''_i , and Y''_i , respectively (recall that the cardinality of each such vertex set is at least w'/c_Δ). We also denote the corresponding set \mathcal{Q}' of paths by \mathcal{Q}'_i . One difficulty is that we are not guaranteed that X''_i is disjoint from $A''_i \cup B''_i$. But it is easy to verify that we can select subsets $A'_i \subseteq A''_i, B'_i \subseteq B''_i, X'_i \subseteq X''_i$ of cardinalities $\lceil w'/(3c_\Delta) \rceil$ each, such that all three sets A'_i, B'_i, X'_i are disjoint. We let $\mathcal{Q}_i \subseteq \mathcal{Q}'_i$ be the set of paths originating at the vertices of X'_i , and we let Y'_i be the set of their endpoints that belong to S_i .

Notice that for each odd integer $1 \leq i \leq \ell'$, we have now selected two subsets $A'_i \subseteq A_i$ and $B'_i \subseteq B_i$ of $\lceil w'/(3c_\Delta) \rceil$ vertices. Using Claim 2.2, we can construct a new Path-of-Sets System $\mathbb{P}' = (\mathcal{C}', \{\mathcal{P}'_i\}_{i=1}^{\lceil \ell'/2 \rceil}, A'_1, B'_{\lceil \ell'/2 \rceil})$ of length $\lceil \ell'/2 \rceil$ and width $\lceil w'/(3c_\Delta) \rceil$, such that $\mathcal{C}' = (C_1, C_3, \dots, C_{2\lceil \ell'/2 \rceil - 1})$, and for each $1 \leq i < \lceil \ell'/2 \rceil$, the paths in \mathcal{P}'_i connect the vertices of B'_{2i-1} to the vertices of A'_{2i+1} .

For each $1 \leq i \leq \lceil \ell'/2 \rceil$, we replace cluster C_{2i-1} with $C'_{2i-1}, \mathcal{Q}_{2i-1}$, and S_{2i-1} , to obtain a hairy Path-of-Sets System of length $\ell'/2$ and width $\lceil w'/(3c_\Delta) \rceil$. \square

B Proof of Lemma 4.6

The proof presented here is almost identical to the proof of [RS83]. We construct a directed graph G' over the set $V(\hat{G})$ of vertices, and prove that this graph is a directed acyclic graph. We will then use the natural ordering of the vertices of $V(G')$ in order to define the bijection μ .

We now define the graph $G' = (V', E')$. The set of vertices of G' is $V' = V(\hat{G})$. The set of edges consists of two subsets. First, for every path $R \in \hat{\mathcal{R}}$, for every pair (v, v') of distinct vertices of R , such that v lies before v' on R , we add a directed edge (v, v') to G' . The resulting edges are called type-1 edges, and the set of these edges is denoted by E_1 . We now define the second set E_2 of edges. Let R, R' be two **distinct** paths of $\hat{\mathcal{R}}$, and let $v \in V(R), v' \in V(R')$ be a pair of vertices. We add a directed edge (v, v') to E_2 iff there is a vertex $u \in V(R)$, that lies strictly after v on R , and there is an edge (u, v') in \hat{G} . The resulting edge (v, v') is called a type-2 edge, and we say that u is the *witness* for this edge. We then let $E' = E_1 \cup E_2$, completing the definition of the graph G' . The following observation is immediate from the definition of the edges of G' .

Observation B.1 *Let $R \in \hat{\mathcal{R}}$ be a path, and let v, v' be two vertices of R , such that v lies strictly before v' on R . Let $u \in V'$ be any vertex, such that $(v', u) \in E'$. Then $(v, u) \in E'$.*

Proof: If u also lies on R , then u appears after v' on R , and hence it also appears after v on R , so $(v, u) \in E_1$. Otherwise, if w is the witness for the edge (v', u) , then w appears after v' on R , and hence $(v, u) \in E_2$, with the witness w . \square

Claim B.2 Graph G' is a directed acyclic graph.

Proof: Assume for contradiction that G' is not acyclic, and let $C = (v_1, v_2, \dots, v_r)$ be a directed cycle in G' , containing fewest vertices (we allow a cycle to consist of two vertices).

Observation B.3 The vertices v_1, \dots, v_r lie on distinct paths of $\hat{\mathcal{R}}$.

Proof: Assume otherwise, so there are indices $1 \leq i < j \leq r$, such that v_i and v_j lie on the same path of $\hat{\mathcal{R}}$. Denote this path by R . In this proof, the addition in the subscripts is modulo r , that is, for $z = r$, $v_{z+1} = v_1$, and for $z = 1$, $v_{z-1} = v_r$.

Assume first that v_i appears before v_j on the path R . Then, from Observation B.1, edge (v_i, v_{j+1}) belongs to E' . Therefore, $v_1, \dots, v_i, v_{j+1}, \dots, v_r$ is a cycle in G' , containing fewer vertices than C (if $j = r$, then we use the cycle (v_1, \dots, v_i)), contradicting the minimality of the cycle C .

Notice that it is impossible that $v_i = v_j$: since graph G' does not contain loops, if $v_i = v_j$, then i and j cannot be consecutive indices, and so $(v_1, \dots, v_i, v_{j+1}, \dots, v_r)$ is a cycle in G' containing fewer vertices than C , a contradiction.

Therefore, v_i appears strictly after v_j on R . In this case, (v_i, v_j) cannot be an edge of G' , and so $j \neq i + 1$. But then, from Observation B.1, edge (v_j, v_{i+1}) belongs to E' , and so (v_{i+1}, \dots, v_j) is a cycle in G' , containing fewer vertices than C , contradicting the minimality of C . \square

We conclude that every vertex of C appears on a distinct path of $\hat{\mathcal{R}}$. For each $1 \leq i \leq r$, let R_i be the path of $\hat{\mathcal{R}}$ containing the vertex v_i , and let u_i be the witness for the edge (v_i, v_{i+1}) . We let $\mathcal{R}' = \{R_1, \dots, R_r\}$ and $\mathcal{R}'' = \hat{\mathcal{R}} \setminus \mathcal{R}'$. For each $1 \leq i \leq r$, let a_i and b_i be the first and the last endpoints of the path R_i , respectively, and let $A' = \{a_1, \dots, a_r\}$, and $B' = \{b_1, \dots, b_r\}$. We construct an (A', B') -linkage $\tilde{\mathcal{R}}$ in \hat{G} , that is different from \mathcal{R}' , and disjoint from \mathcal{R}'' . This implies that $\tilde{\mathcal{R}} \cup \mathcal{R}''$ is an (\hat{A}, \hat{B}) -linkage in \hat{G} , that is different from $\hat{\mathcal{R}}$, contradicting the uniqueness of $\hat{\mathcal{R}}$.

The linkage $\tilde{\mathcal{R}}$ consists of r paths R'_1, \dots, R'_r , where for all $1 \leq i \leq r$, path R'_i starts at a_i and terminates at b_{i-1} (path R'_1 starts at a_1 and terminates at b_r).

For $1 \leq i \leq r$, path R'_i consists of: (i) the segment of R_i from a_i to v_i ; (ii) the edge (v_i, u_{i-1}) of \hat{G} (for $i = 1$, we use edge (v_1, u_r)); and (iii) the segment of R_{i-1} from u_{i-1} to b_{i-1} (see Figure 5).

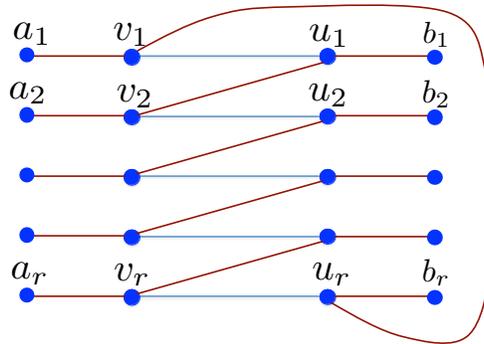


Figure 5: Rerouting the paths in \mathcal{R}' . The new set $\tilde{\mathcal{R}}$ of paths is shown in red.

\square

For convenience, let $\hat{n} = |V(\hat{G})|$. We now define the bijection μ of the vertices of $V(\hat{G})$ to $\{1, \dots, \hat{n}\}$, in a natural way, using G' . We perform \hat{n} iterations, where in iteration i we select a vertex $v_i \in V(G')$, setting $\mu(v_i) = i$. The first iteration is executed as follows. Since G' is a directed acyclic graph, there is some vertex $v \in V'$ with no incoming edges. We denote $v_1 = v$, and set $\mu(v) = 1$. Consider now some iteration $i > 1$, and assume that vertices v_1, \dots, v_{i-1} were already defined. Consider the graph $G'_i = G' \setminus \{v_1, \dots, v_{i-1}\}$. This graph is again a directed acyclic graph, so it contains some vertex v with no incoming edges. We let $v_i = v$, and we set $\mu(v_i) = i$. This finishes the definition of the mapping μ . We now verify that it has the required properties.

For the first property, if $R \in \hat{\mathcal{R}}$ is any path, and v, v' are distinct vertices of R , with v' appearing before v on R , then the edge (v', v) is present in G' , and so $\mu(v') < \mu(v)$.

Consider now some integer $1 \leq t \leq \hat{n}$. Recall that set S_t contains, for every path $R \in \hat{\mathcal{R}}$, the first vertex v on R with $\mu(v) \geq t$; if no such vertex exists, then we add the last vertex of R to S_t . Recall also that $Y_t = \{v \in V(\hat{G}) \mid \mu(v) < t\}$ and $Z_t = \{v \in V(\hat{G}) \mid \mu(v) \geq t\}$. We need to show that graph $\hat{G} \setminus S_t$ contains no path connecting a vertex of Y_t to a vertex of Z_t . Assume otherwise. Then there must be some pair of vertices $v \in Y_t \setminus S_t$ and $v' \in Z_t \setminus S_t$, such that $(v, v') \in E(\hat{G})$. Let $R \in \hat{\mathcal{R}}$ be the path containing v , and let $R' \in \hat{\mathcal{R}}$ be the path containing v' .

Assume first that $R = R'$. Since $v \in Y_t \setminus S_t$, $\mu(v) < t$, and, since $v' \in Z_t \setminus S_t$, $\mu(v') \geq t$. Moreover, there must be some other vertex u of R that belongs to S_t , and, from the definition of S_t , u appears between v and v' on R . Let a, b be the endpoints of the path R . Then we can construct a new path R' , connecting a to b , as follows: path R' consists of the segment of R from a to v ; the edge (v, v') , and the segment of R from v' to b . But then, replacing R with R' in $\hat{\mathcal{R}}$, we obtain an (\hat{A}, \hat{B}) -linkage in \hat{G} that is different from $\hat{\mathcal{R}}$, contradicting the unique linkage property of \hat{G} .

Therefore, $R \neq R'$ must hold. Let x be the vertex of R' that belongs to S_t . From the definition of S_t , $\mu(x) \geq t$, and x appears before v' on path R' . But then there must be an edge (x, v) in $E_2 \subseteq E(G')$, whose witness is v' . From the definition of μ , if such an edge is present in G' , then $\mu(x) < \mu(v)$ must hold. But $\mu(x) \geq t$, while $\mu(v) < t$, a contradiction.