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

Julia Chuzhoy† Zihan Tan‡

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 $f : \mathbb{Z}^+ \to \mathbb{Z}^+$, 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 $f(g) = O(g^{98} \operatorname{poly}\log g)$ was proved by Chekuri and Chuzhoy. It was later improved to $f(g) = O(g^{36} \operatorname{poly}\log g)$, and then to $f(g) = O(g^{19} \operatorname{poly}\log g)$. In this paper we further improve this bound to $f(g) = O(g^9 \operatorname{poly}\log g)$. 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 $f(g)$.

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 $f : \mathbb{Z}^+ \to \mathbb{Z}^+$, 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 $f(g) = \Omega(g^2 \log g)$ using constant-degree $\Omega(\log n)$-girth 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 $f(g) = O(g^{98} \operatorname{poly}\log g)$ 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 $f(g) = O(g^{36} \operatorname{poly}\log g)$, 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 $f(g) = O(g^{19} \operatorname{poly}\log g)$ 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*

*Supported in part by NSF grants CCF-1318242 and CCF-1616584.

†Toyota Technological Institute at Chicago. Email: cjulia@ttic.edu.

‡Computer Science Department, University of Chicago. Email: zihantan@uchicago.edu.

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 improving the best current upper bounds for the Excluded Grid Theorem, we believe that our framework is significantly simpler than 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. Moreover, there are natural barriers that we discuss below, which 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 it.

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 the notion 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 length ℓ and width w (see Figure 1(a)) consists of a sequence $\mathcal{C} = (C_1, \ldots, C_\ell)$ of ℓ connected sub-graphs of the input graph G that we call *clusters*. For each cluster $C_i \subseteq V(G)$, we are given two disjoint subsets $A_i, B_i \subseteq C_i$ of its vertices of cardinality w each. We require that the vertices of $A_i \cup B_i$ are well-linked in $G[C_i]$[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 length g^2 and width g^2, then G contains the $(\Omega(g) \times \Omega(g))$-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 both length and width $\Omega(g^2)$.

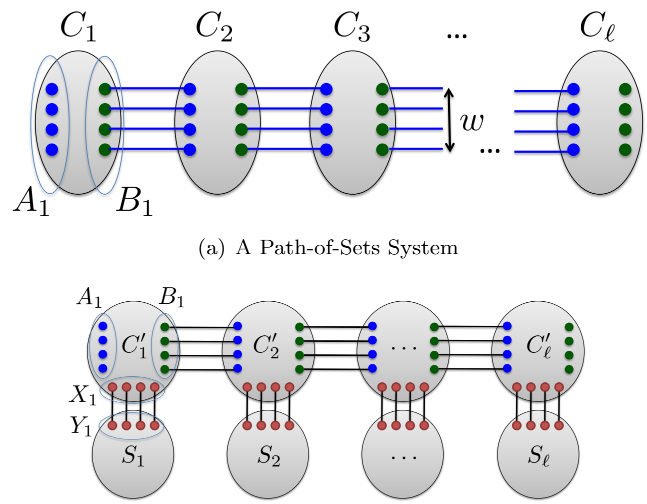

(a) A Path-of-Sets System

(b) A hairy Path-of-Sets System

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 subgraph 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-linked[2] 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. While the bound on $f(g)$ that this result produces is weak: $f(g) = O(g^{98} \operatorname{poly}\log g)$, 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

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

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

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 can obtain a Path-of-Sets System of length ℓ and width $\Omega(k/\operatorname{poly}\log k)$. This fact was used in [CC15] to show that any graph G of treewidth k contains a sub-graph G' of treewidth $\Omega(k/\operatorname{poly}\log k)$, 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 on the Excluded Grid Theorem.

In [Chu15], a different strategy for obtaining the 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 length 1 and width $\Omega(k)$. 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 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, we employ a subroutine, 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, we obtain 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 in the final bound on $f(g)$ that we obtain, 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 $w = \Omega(k/\operatorname{poly}\log k)$, 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 $C'_1, \ldots, C'_\ell$ 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 our input graph G contains the $(g \times g)$-grid as a minor. Intuitively, a crossbar inside cluster C_i of a Path-of-Sets System consists of a set $\mathcal{P}^*_i$ of g^2 node-disjoint paths, connecting vertices of A_i to vertices of B_i; and another set $\mathcal{Q}^*_i$ of g^2 node-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 structural result that we prove is that, if the width w' of the hairy Path-of-Sets System is sufficiently large, then for each cluster C_i, either it contains a crossbar; or it 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 input 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 for routing on expanders to show that such an expander must contain the $(g \times g)$-grid as a minor.

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. The proof of the structural theorem appears in Section 4.

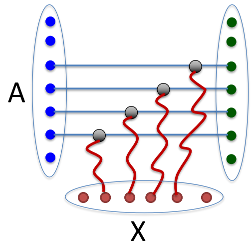

Figure 2: A Crossbar. The paths of $\mathcal{P}^*$ are shown in blue and the paths of $\mathcal{Q}^*$ in red.

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 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. We use standard notions of well-linkedness and linkedness.

Definition 2.1. 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 any two disjoint subsets T', T'' 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 any two disjoint subsets T', T'' 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 some 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 2.2. 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 any two subsets $A' \subseteq A, B' \subseteq B$ of vertices, there is a set $\mathcal{P}$ of node-disjoint paths in G connecting vertices of A' to vertices of B', with $|\mathcal{P}| = \min\{|A'|, |B'|\}$.

[3]Note 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.

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 2.3. 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 $\mathcal{C} = (C_1, \ldots, C_\ell)$ 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 from a vertex of $B_i \subseteq C_i$, terminates at a vertex of $A_{i+1} \subseteq C_{i+1}$, 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 $g' = \lfloor g/c \rfloor$.*

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 $\mathcal{C} = (C_1, \ldots, C_\ell)$. 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 full version of the paper.

CLAIM 2.1. *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 $\hat{\mathbb{P}} = (\hat{\mathcal{C}}, \left\{\hat{\mathcal{P}}_i\right\}_{i=1}^{\lceil \ell/2 \rceil - 1}, \hat{A}_1, \hat{B}_{\lceil \ell/2 \rceil})$ of length $\lceil \ell/2 \rceil$ and width w', such that $\hat{\mathcal{C}} = (C_1, C_3, \ldots, C_{2\lceil \ell/2 \rceil - 1})$; 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 2.1, 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 2.4. 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 $\mathcal{S} = (S_1, \ldots, S_\ell)$ 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 will sometimes 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 $k/\operatorname{poly}\log k = \Omega(w\ell^{48})$, 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 $w = k/\operatorname{poly}\log k$.

THEOREM 2.2. *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 the proof of Theorem 1.1, with some details delayed to Section 4. We start by introducing the main new combinatorial object that we use, called a *crossbar*.

DEFINITION 3.1. 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 $\kappa \geq 2^{22} g^9 \log g$. 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 Section 4. The following theorem will be used 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 $\mathbb{H} = (\mathcal{C}, \mathcal{S}, \{\mathcal{P}_i\}_{i=1}^{\ell-1}, \{\mathcal{Q}\}_{i=1}^{\ell}, A_1, B_\ell)$ 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, X_i)-crossbar in C_i of width g^2. Then G contains the $(g' \times g')$-grid as a minor, for $g' = \Omega(g/\log^5 g)$.*

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 $(\Omega(g/\operatorname{poly}\log g) \times (\Omega(g/\operatorname{poly}\log g))$.

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 2.2, hold for ℓ, w and k. From Theorem 2.2, 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 cluster C_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 $\tilde{\mathcal{Q}}$ of w node-disjoint paths connecting vertices of A_i to vertices of X_i in C_i. We can therefore apply Theorem 3.1 to the cluster C_i, and the sets A_i, B_i, X_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, 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, X_i)-crossbar in C_i of width g^2. But then from Theorem 3.2, G contains the $(g' \times g')$-grid as a minor, for $g' = \Omega(g/\log^5 g)$. 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 at least 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 at least k, then it contains a grid minor of size $(g' \times g')$.

One last issue is that we need to replace the $\operatorname{poly}\log k$ in the above bound on k by $\operatorname{poly}\log g$, 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, $k'(g) \leq cg^{10}$, and so $\log k'(g) = O(\log g)$. 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.

We now provide a high-level sketch of the proof of Theorem 3.2; a formal proof is deferred to the full version of the paper. For simplicity, assume that ℓ is an even integer. For every odd integer $1 \leq i \leq \ell$, let $(\mathcal{P}^*_i, \mathcal{Q}^*_i)$ be the (A_i, B_i, X_i)-crossbar of width g^2 in C_i, and let $A^*_i \subseteq A_i$, $B^*_i \subseteq B_i$ be the sets of endpoints of the paths in $\mathcal{P}^*_i$, lying in A_i and B_i, respectively, so that $|A^*_i| = |B^*_i| = g^2$. Using the stitching claim (Claim 2.1), we can obtain a hairy Path-of-Sets System $\mathbb{H}' = (\mathcal{C}', \mathcal{S}', \{\mathcal{P}'_i\}_{i=1}^{\ell/2-1}, \{\mathcal{Q}'_i\}_{i=1}^{\ell/2}, A'_1, B'_\ell)$ of length $\ell/2 = \Theta(\log g)$ and width g^2, where $\mathcal{C}' = (C_1, C_3, \ldots, C_{\ell-1})$, $\mathcal{S}' = (S_1, S_3, \ldots, S_{\ell-1})$, and for all $1 \leq i < \ell/2$, set $\mathcal{P}'_i$ connects vertices of $B'_i = B^*_{2i-1}$ to vertices of $A'_{i+1} = A^*_{2i+1}$. Moreover, $A'_1 = A^*_1$ and $B'_{\ell/2} = B^*_{\ell-1}$.

For convenience, for each $1 \leq i \leq \ell/2$, we rename the crossbar contained in C_{2i-1} by $(\mathcal{P}^*_i, \mathcal{Q}^*_i)$. Recall that we denoted the corresponding sets of endpoints of the paths in $\mathcal{P}^*_i$ as A'_i and B'_i. By concatenating the paths in $\mathcal{P}^*_1, \mathcal{P}_1, \ldots, \mathcal{P}_{\ell/2-1}, \mathcal{P}^*_{\ell/2}$, we obtain a collection $\mathcal{R}$ of paths, such that for each $R \in \mathcal{R}$, for each $1 \leq i \leq \ell/2$, $R \cap C_{2i-1}$ is a path of $\mathcal{P}^*_i$. It is now easy to show, using the Cut-Matching Game of [KRV09], that there is an expander graph H on g^2 vertices and maximum vertex degree $O(\log g)$, such that H is a minor of G. Moreover, we can construct a model φ of H in G, such that for every vertex $v \in V(H)$, $\varphi(v)$ is a path of $\mathcal{R}$. 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). Recently, Chuzhoy and Nimavat [CN18] showed that any expander H on n vertices, with maximum vertex degree d, contains any graph H' with at most $n/(d^c \log n)$ edges and vertices as a minor, for some constant c. We could use their result as a black-box, but, as their result tries to optimize the bounds they obtain on the minor size, its proof is more involved than what we need here. Instead, we provide a direct proof, that exploits the fact that routing on expander graphs is easy. In particular, it is known that, for any partition $U_1, \ldots, U_{2r}$ of the vertices of the expander into large enough subsets ($|U_i| = \operatorname{poly}\log n$ for all i is sufficient), and any ordering of these subsets, there is a set $\{P_1, \ldots, P_r\}$ of node-disjoint paths, where for each $1 \leq i \leq r$, path P_i connects some vertex of U_{2i-1} to some vertex of U_{2i}. We use all except the first cluster C_1 of the hairy Path-of-Sets System in order to embed an expander H over g^2 vertices into G. Recall that the vertices of the expander are embedded into the paths of $\mathcal{R}$. We then use the first cluster in order to "group" these paths into groups of a large enough cardinality, such that the paths participating in every group can be connected to each other inside C_1. The subgraphs of C_1 spanning these groups become the embeddings of the vertices of the grid. We use the above mentioned result about routing in expanders in order to embed the edges of the grid minor.

4 Building the Crossbar

This section is dedicated to 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 $\kappa \geq 2^{22} g^9 \log g$. 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 $\Omega(g^2)$.

As our first step, we 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 $\mathcal{R}_1, \ldots, \mathcal{R}_D$ of disjoint subsets of paths in $\mathcal{P}$ (that is, for all $1 \le i \le D$, $\mathcal{R}_i \subseteq \mathcal{P}$), such that for all i, $|\mathcal{R}_i| \le g^2$. 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 $1 \le i \le D$, 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 $\sigma_1(R), \ldots, \sigma_M(R)$ 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 $1 \le i \le M$, 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 $1 \le i \le M$, 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 $1 \le i \le M$, 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 note that the above high-level exposition of the proof can be exploited to obtain slightly weaker bounds for Theorem 3.1, that is, we need to assume that $\kappa = \Omega(g^{10})$, leading to a weaker bound of $f(g) = O(g^{10} \operatorname{poly}\log g)$ for Theorem 1.1. A slightly more involved process that is formally described below gives a full proof of Theorem 3.1 with the claimed bounds.

We now provide a formal proof of Theorem 3.1, 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 4.1. Let $D > 0$ be an integer. A pseudo-grid of depth D consists of the following two ingredients. The first ingredient is a family $\{\mathcal{R}_1, \mathcal{R}_2, \ldots, \mathcal{R}_D\}$ of subsets of $\mathcal{P}$, where for all $1 \le i \le D$, $|\mathcal{R}_i| \le g^2$, and for all $1 \le i \ne j \le D$, $\mathcal{R}_i \cap \mathcal{R}_j = \emptyset$. 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, $|\mathcal{P}'| \ge |\mathcal{Q}'| = \lceil \kappa/4 \rceil$). 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 $1 \le i \le D$, 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 $1 \le D \le \kappa/(4g^2)$. Then either H contains an (A, B, X)-crossbar of width g^2, or it contains a pseudo-grid of depth D.*

Proof. Assume first that at least $\kappa/2$ paths of $\mathcal{P}$ contain vertices of X. In this case, we can obtain an (A, B, X)-crossbar of width g^2 as follows. We let $\mathcal{P}^*$ contain all paths $P \in \mathcal{P}$ with $P \cap X \ne \emptyset$. As long as $|\mathcal{P}^*| > g^2$, we discard paths from $\mathcal{P}^*$ arbitrarily, until $|\mathcal{P}^*| = g^2$ holds. For each path $P \in \mathcal{P}^*$, the corresponding path Q_P consists of a single vertex in $P \cap X$. It is easy to verify that we obtain an (A, B, X)-crossbar of width g^2. Therefore, we assume from now on that at least $\kappa/2$ paths of $\mathcal{P}$ are disjoint from X, and we denote the set of all such paths by $\mathcal{P}'_0$, so $|\mathcal{P}'_0| \ge \kappa/2$.

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 $1 \le i \le D$, we will denote by $\mathcal{P}'_i = \mathcal{P} \setminus (\mathcal{R}_1 \cup \cdots \cup \mathcal{R}_i)$ the collection of the remaining paths of $\mathcal{P}$. We will ensure that for all i, $|\mathcal{R}_i| \le g^2$.

We now describe the ith iteration of our algorithm, whose input is a set $\mathcal{P}'_{i-1} \subseteq \mathcal{P}$ of at least $\kappa/2 - (i - 1)g^2$ paths. In order to execute the ith 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 $S_i = \left\{v_P \mid P \in \mathcal{P}'_{i-1}\right\}$ 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 $\mathcal{Q}^* = \left\{Q' \mid Q \in \hat{\mathcal{Q}}\right\}$. 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 $\mathcal{Q}^* = \left\{Q' \mid Q \in \hat{\mathcal{Q}}\right\}$, 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 $\mathcal{P}'_i = \mathcal{P} \setminus (\mathcal{R}_1 \cup \cdots \cup \mathcal{R}_i) = \mathcal{P}'_{i-1} \setminus \mathcal{R}_i$. 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 4.1. *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, obtaining the sets $\mathcal{R}_1, \ldots, \mathcal{R}_D \subseteq \mathcal{P}$ 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 $D \leq \kappa/(4g^2)$, 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 $\{\mathcal{R}_1, \ldots, \mathcal{R}_D\}$ 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 Observation 4.1, 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 4.1, $\sigma_i(Q)$ 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 $|\mathcal{R}| = N$, where $D \leq N \leq Dg^2$. 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. The proof is deferred to the full version.

Observation 4.2. *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 $N-1$.*

We need the following definition.

Definition 4.2. 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 (I1). 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 v, by unifying v with one of its neighbors; update the corresponding path of $\mathcal{Q}''$.

Let H'' be the graph obtained at the end of this procedure. The proof of the following simple observation is deferred to the full version of the paper.

Observation 4.3. *Graph H'' is a minor of H and Property (I1) holds in H''. Moreover, H'' has the perfect unique linkage property with respect to A' and B', with the unique linkage being $\mathcal{R}$.*

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 $\mathcal{R}$. However, we may need to perform this slicing twice, with two different sets of parameters, so the definition that we give here, and the following theorem asserting the existence of the slicing are stated in general terms.

Definition 4.3. 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 $\Lambda(R) = (v_0(R), v_1(R), \ldots, v_{\hat{M}}(R))$ 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 $v_0(R), v_1(R), \ldots, v_{\hat{M}}(R)$ 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 $\Sigma_i = \left\{\sigma_i(R) \mid R \in \hat{\mathcal{R}}\right\}$, 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 4.2. *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 Lemma 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.

Lemma 4.1. *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 $\mu : V(\hat{G}) \to \left\{1, \ldots, |V(\hat{G})|\right\}$ 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 $Y_t = \left\{v \in V(\hat{G}) \mid \mu(v) < t\right\}$ and $Z_t = \left\{v \in V(\hat{G}) \mid \mu(v) \geq t\right\}$. Then μ has the following properties: (i) 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 (ii) 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 4.1 to graph $\hat{G}$ and the sets $\hat{A}$ and $\hat{B}$ of its vertices, obtaining a bijection $\mu : V(\hat{G}) \to \left\{1, \ldots, |V(\hat{G})|\right\}$.

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 $\left\{R_1(S_t) \mid R \in \hat{\mathcal{R}}\right\}$, and let $\mathcal{Q}^2(S_t)$ be defined similarly for $\left\{R_2(S_t) \mid R \in \hat{\mathcal{R}}\right\}$. 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. The proof is deferred to the full version of the paper.

Observation 4.4. *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 $1 \leq t < |V(\hat{G})|$, $|\mathcal{Q}^1(S_{t+1}) \setminus \mathcal{Q}^1(S_t)| \leq 1$.*

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: (i) $|\hat{\mathcal{Q}}_{i+1}| \geq |\hat{\mathcal{Q}}| - (i+2)\hat{N} - i\hat{w}$; and (ii) for each path $R \in \mathcal{R}$, 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 4.4. 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 $\left\{v_i(R) \mid R \in \hat{\mathcal{R}}\right\} = S_{t_i}$, 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 4.4, 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 $v_0(R), \ldots, v_i(R)$ remain the same as before. We let $v_{i+1}(R)$ be the unique vertex of $R \cap S_{t+1}$, 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 $R \in \hat{\mathcal{R}}$, $\Lambda'(R) = (v_0(R), \ldots, v_{i+1}(R))$. 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 $\hat{N}$ of them. 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 $M_1 = 128g^3 \log g$. From Theorem 4.2, we can obtain an M_1-slicing $\Lambda = \{\Lambda(R)\}_{R \in \mathcal{R}}$ of $\mathcal{R}$, of width $w = 2^{11}g^6$ with respect to $\mathcal{Q}''$, since $N \leq g^2 D = 64g^6$, and so $|\mathcal{Q}''| \geq \kappa/8 \geq 2^{19}g^9 \log g \geq M_1(2N + w)$. 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 *ith slice of* Λ.

4.3 Step 3: Intersecting Path Sets.

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

Definition 4.4. 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 4.2. *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, $|\hat{\mathcal{Q}} \setminus \hat{\mathcal{Q}}'| \leq |\hat{\mathcal{R}}|\hat{w}/\hat{D} \leq \hat{\mathcal{Q}}/2$. □

For each $1 \leq i \leq M_1$, We apply Lemma 4.2 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 $\mathcal{Q}_i'$.

We now distinguish between two cases. For $1 \leq i < M_1$, 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.

If Case 1 happens, then we proceed directly to Step 4. Assume now that Case 2 happens, and consider some type-2 slice $(\Sigma_i, \mathcal{Q}_i)$. 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 4.3. *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 4.3. *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 $\hat{M} = g$.*

Recall that $M_1 = 128g^3 \log g$, and so at least $64g^3 \log g$ of the slices of Λ are type-2 slice. Each such slice $(\Sigma_i, \mathcal{Q}_i)$ is in turn sliced into g slices. 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 4.3. 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 $v_1(\sigma) = v_2(\sigma) = \cdots = v_{\hat{M}}(\sigma)$ 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 4.1. *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 jth 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. 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 $\tilde{\mathcal{Q}}_i$. The proof of the next observation is deferred to the full version of the paper.

Observation 4.5. *Graph H_i' has the perfect unique linkage property with respect to (A_i', B_i'), with the unique linkage being Σ_i'.*

Recall that $|\Sigma_i'| \leq N/g$, and $|\tilde{\mathcal{Q}}_i| \geq 512g^6$. We now invoke Theorem 4.2 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 $v_1(\sigma) = v_2(\sigma) = \cdots = v_{\hat{M}}(\sigma)$ 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. $\square$

If Theorem 4.3 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 ith slice of this new slicing by $(\Sigma_i, \mathcal{Q}_i)$, where $\mathcal{Q}_i \subseteq \mathcal{Q}''$ and Σ_i contains the ith segment of each path $R \in \mathcal{R}$.

As in Case 1, for each $1 \leq i \leq M_2$, we employ Lemma 4.2 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. In order to do so, we need to verify that $|\mathcal{Q}_i| \geq 2|\Sigma_i| \cdot 4g^2/(D/2) = 16|\Sigma_i|g^2/D$. Recall that $|\mathcal{Q}_i| \geq N/g$, while $|\Sigma_i| \leq N$, and $D = 64g^4$, so the inequality indeed holds. As before, every path of $\mathcal{Q}_i$ intersects at least D paths of Σ_i, and so the conditions of Lemma 4.2 are satisfied, with $\hat{w} = 4g^2$ and $\hat{D} = D/2$.

Summary of Step 3. In Case 1, at the end of Step 3 we obtain an M_1-slicing of $\mathcal{R}$, with $M_1 = 128g^3 \log g$, whose width with respect to $\mathcal{Q}''$ is $w_1 = 2^{11}g^6$. In Case 2, we obtain an M_2-slicing of $\mathcal{R}$, with $M_2 = 8g^4 \log g$, whose width with respect to $\mathcal{Q}''$ is $w_2 = N/g$. In either case, for each slice $(\Sigma_i, \mathcal{Q}_i)$, we also computed a partition (Σ_i', Σ_i'') of Σ_i, and a subset $\mathcal{Q}_i' \subseteq \mathcal{Q}_i$ of paths, such that $\mathcal{Q}_i' \neq \emptyset$, and $(\Sigma_i', \mathcal{Q}_i')$ are $(4g^2, D/2)$-intersecting. Since $\mathcal{Q}_i' \neq \emptyset$, we are guaranteed that $|\Sigma_i'| \geq D/2$. Moreover, in Case 1, for at least half the indices i, $(\Sigma_i, \mathcal{Q}_i)$ is a type-1 slice, that is, $|\Sigma_i'| \geq N/g$. If Case 1 happens, 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 $(\Sigma_i, \mathcal{Q}_i)$, $|\Sigma_i'| \geq N/g$.

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 4.5. 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 4.6. *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, whose proof is deferred to the full version.

OBSERVATION 4.7. *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 $|E(X,Y)| < \min\{\hat{w}, |T \cap X|, |T \cap Y|\}$.*

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 4.4. *Let $\hat{G}$ be a graph, and let $\hat{\Sigma}$ and $\hat{\mathcal{Q}}$ be two sets of node-disjoint paths in $\hat{G}$ (but a path in $\hat{\Sigma}$ and a path in $\hat{\mathcal{Q}}$ may intersect). Let $\hat{w}, \hat{D} > 0$ be integers, such that $(\hat{\Sigma}, \hat{\mathcal{Q}})$ 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: (i) 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); (ii) $|\hat{\Sigma}'| \geq |\hat{\Sigma}|/4$; and (iii) every path $\sigma \in \hat{\Sigma}'$ belongs to some 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 4.7, 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 $\mathrm{out}(C)$ 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 $|\mathrm{out}(C)| < 4\hat{w}$, and $\mathcal{C}_2$ contains all remaining clusters of $\mathcal{C}$. The proofs of the following two observations are deferred to the full version of the paper.

OBSERVATION 4.8. $|\mathcal{C}_1| \geq r/2$.

OBSERVATION 4.9. *Every cluster $C \in \mathcal{C}_1$ is happy.*

COROLLARY 4.1. $r \leq 2|\hat{\Sigma}|/\hat{D}$.

Proof. Note that from Observations 4.8 and 4.9, $\sum_{C \in \mathcal{C}_1} |\hat{\Sigma}(C)| \geq r\hat{D}/2$. On the other hand, $\sum_{C \in \mathcal{C}_1} |\hat{\Sigma}(C)| \leq |\hat{\Sigma}|$. The corollary now follows.

□

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, $|\hat{\Sigma}'| \geq |\hat{\Sigma}|/4$. □

Assume first that Case 2 happened, so we have $M_2 = 8g^4 \log g$ slices $\{(\Sigma_i, \mathcal{Q}_i)\}_{1\le i\le M_2}$. Recall that for each $1 \le i \le M_2$, we have computed subsets $\Sigma'_i \subseteq \Sigma_i$ and $\mathcal{Q}'_i \subseteq \mathcal{Q}_i$, such that $(\Sigma'_i, \mathcal{Q}'_i)$ are $(4g^2, D/2)$-intersecting. Let H'_i be a graph obtained from the union of the paths in Σ_i and $\mathcal{Q}_i$. 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. Note that $\hat{D} \ge 8\hat{w}$, since $D = 64g^4$. We apply Theorem 4.4 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, and let $C_i \in \mathcal{C}_i$ be any such cluster. We denote by $\tilde{\Sigma}_i = \hat{\Sigma}(C_i)$. Recall that $|\tilde{\Sigma}_i| \ge D/4$, and the endpoints of the paths in $\tilde{\Sigma}_i$ are g^2-weakly well-linked in C_i. Finally, we let $\tilde{\mathcal{C}} = \{C_1, \ldots, C_{M_2}\}$. To summarize, we have obtained a collection of M_2 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.

Assume now that Case 1 happened, so we have $M_1 = 64g^3 \log g$ slices $\{(\Sigma_i, \mathcal{Q}_i)\}_{1\le i\le M_1}$. As before, for each $1 \le i \le M_1$, we have computed subsets $\Sigma'_i \subseteq \Sigma_i$ and $\mathcal{Q}'_i \subseteq \mathcal{Q}_i$, such that $(\Sigma'_i, \mathcal{Q}'_i)$ are $(4g^2, D/2)$-intersecting. But now we are guaranteed that $|\Sigma'_i| \ge N/g$. As before, let H'_i be a graph obtained from the union of the paths in Σ_i and $\mathcal{Q}_i$. As before, 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 $\hat{D} \ge 8\hat{w}$.

As before, for each $1 \le i \le M_1$, we apply Theorem 4.4 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 $|\tilde{\Sigma}(C)| \ge D/4$. We denote $\tilde{\Sigma}_i = \bigcup_{C\in\mathcal{C}_i} \tilde{\Sigma}(C_i)$. Notice that Theorem 4.4 guarantees that $|\tilde{\Sigma}_i| \ge |\Sigma'_i|/4 \ge N/(4g)$. Let $\mathcal{C} = \bigcup_{i=1}^{M_1} \mathcal{C}_i$, and let $\tilde{\Sigma} = \bigcup_{i=1}^{M_1} \tilde{\Sigma}_i$, so $|\tilde{\Sigma}| \ge M_1 N/4g$.

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 $C \in \mathcal{C}$, $D/4 \le |\tilde{\Sigma}(C)| \le N \le g^2 D$. For $0 \le j < 2\log g + 2$, we say that cluster C belongs to class $\mathcal{S}_j$ iff $D \cdot 2^j/4 \le |\tilde{\Sigma}(C)| < D \cdot 2^{j+1}/4$. 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 $\frac{|\tilde{\Sigma}|}{2\log g+2} \ge \frac{M_1 N}{16 g \log g}$. 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 $|\tilde{\mathcal{C}}'| \ge \frac{|\tilde{\Sigma}'|}{D\cdot 2^{j+1}/4} \ge \frac{M_1 N}{4D\cdot 2^{j+1} g\log g} = \frac{64 g^3 N \log g}{8D\cdot 2^j g \log g} = \frac{8Ng^2}{D\cdot 2^j}$, while for each $C \in \mathcal{C}$, $|\tilde{\Sigma}(C)| \ge D \cdot 2^j/4$.

To summarize, in Case 1, 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 $C \ne C'$, 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 $C \in \tilde{\mathcal{C}}$, $|\tilde{\Sigma}(C)| \ge D \cdot 2^j/4$, and the endpoints of the paths of $\tilde{\Sigma}(C)$ are g^2-weakly well-linked in C.

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 $[\hat{N}] = \left\{1, \ldots, \hat{N}\right\}$. Let $\pi : \mathcal{R} \to [N]$ 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 4.5. *Let $\hat{D}, \hat{N}, \hat{M}, \hat{w}$ be non-negative integers, such that (i) $\hat{N} \ge 3\hat{w}$; (ii) $\hat{D}^2 \ge 4\hat{N}\hat{w}$; and (iii) $\hat{M}\hat{D} \ge 2\hat{N}\hat{w}$ hold, and let $S_1, \cdots, S_{\hat{M}}$ be subsets of $[\hat{N}]$, where for each $1 \le i \le \hat{M}$, $|S_i| \ge \hat{D}$. Then there are $\hat{w}$ indices $1 \le i_1 < i_2 < \cdots < i_{\hat{w}} \le \hat{M}$, such that for all $1 \le j < \hat{w}$, $|S_{i_j} \cap S_{i_{j+1}}| \ge \hat{w}$.*

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. Assume first that Case 2 happens. We then denote $\hat{M} = M_2 = 8g^4 \log g$, $\hat{D} = D/4$, $\hat{N} = N$, and $\hat{w} = g^2$. Recall that we are given a set $\mathcal{C} = \left\{C_1, C_2, \ldots, C_{\hat{M}}\right\}$ of clusters, where cluster C_i corresponds to slice i. Recall also that for each i, $|\tilde{\Sigma}(C_i)| \ge D/4$. We now build the subsets $S_1, \ldots, S_{\hat{M}}$ of $[\hat{N}]$ as follows. Fix some $1 \le i \le \hat{M}$. 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, $|S_i| \ge D/4 = \hat{D}$. We now verify that the conditions of Theorem 4.5 hold for the chosen parameters. The first condition, $\hat{N} \ge 3\hat{w}$, 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: $DM_2 \geq 8Ng^2$. Using the fact that $N \leq Dg^2$, and that $M_2 \geq 8g^4$, the inequality clearly holds.

Therefore, we can now apply Theorem 4.5 to conclude that there are indices $1 \leq i_1 < i_2 < \cdots < i_{g^2} \leq \hat{M}$, 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 $C'_1, \ldots, C'_{g^2}$ 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 $\Sigma^A_{i_j} \subseteq \tilde{\Sigma}_{i_j}$ 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 $\Sigma^B_{i_j} \subseteq \tilde{\Sigma}_{i_j}$ 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 4.4, the endpoints of the paths in $\Sigma^A_{i_j} \cup \Sigma^B_{i_j}$ 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 4.6, $A_j \cup B_j$ 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 $\sigma_{i_j+1}(R), \ldots, \sigma_{i_{j+1}-1}(R)$, and two additional edges, one immediately preceding $\sigma_{i_j+1}(R)$, and one immediately following $\sigma_{i_{j+1}-1}(R)$. Since $i_1 < i_2 < \cdots < i_{g^2}$, the paths in $\bigcup_j \mathcal{P}_j$ are disjoint from each other, and are internally disjoint from $\bigcup_{j'} C'_{j'}$.

Assume now that Case 1 happens. We denote $\hat{M} = |\tilde{\mathcal{C}}| \geq \frac{8Ng^2}{D \cdot 2^j}$, $\hat{w} = g^2$, $\hat{N} = N$, and $\hat{D} = \frac{D \cdot 2^j}{4}$. Recall that for each $C \in \tilde{\mathcal{C}}$, $|\tilde{\Sigma}(C)| \geq \hat{D}$. Again, we need to verify that the conditions of Theorem 4.5 hold for this 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 before, this condition continues to hold. The second condition is $\hat{D}^2 \geq 4\hat{N}\hat{w}$. Recall that this condition held for Case 2, 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.

Recall that set $\mathcal{C}$ may contains 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 $C_1, \ldots, C_{\hat{M}}$ be the resulting ordering of the clusters. We define the sets $S_1, \ldots, S_{\hat{M}}$ 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 now use Theorem 4.5 to obtain a sequence $1 \leq i_1 < i_2 < \cdots < i_{g^2} \leq \hat{M}$ 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 $C_{i_1}, C_{i_2}, \ldots, C_{i_{g^2}}$ must belong to a different slice. The remainder of the construction of the Path-of-Sets system is done exactly as in Case 2.

In order to complete the proof of Theorem 3.1, it is now enough to prove Theorem 4.5.

Proof of Theorem 4.5. The proof consists of three steps. First, we use the sets $S_1, \ldots, S_{\hat{M}}$ 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 desires sequence $i_1, \ldots, i_w$ of indices.

We start by defining a directed graph $\hat{G} = (\hat{V}, \hat{E})$, where $\hat{V} = \{1, 2, \cdots, \hat{M}\}$, 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 use the following claim, whose proof is deferred to the full version of the paper.

Claim 4.2. *Let $\mathcal{S} \subseteq \{S_1, \ldots, S_{\hat{M}}\}$ be any collection of $r = \lceil 2\hat{N}/\hat{D} \rceil$ 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 4.2. *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 $|V'| < 2\hat{N}/\hat{D}$.*

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. The proof of the next claim is deferred to the full version of the paper.

Claim 4.3. *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 $\ell(G) \geq \hat{M}/\alpha(G)$.*

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 4.5). This directed path immediately defines the desired sequence $i_1, \ldots, i_{\hat{w}}$ of indices, completing the proof of Theorem 4.5. □

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 $\mathcal{C} = (C_1, \ldots, C_\ell)$. In this step we complete the proof 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 4.6. 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 4.10. *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 4.6. (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 4.7. (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 $\mathcal{C} = (C_1, \ldots, C_\ell)$, 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 4.6 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 4.6 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 4.6 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 4.7, (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 $\tilde{w} = \Omega(g^2)$.

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.

[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 Erdös–Pósa property for minor-closed graph classes. *Journal of Graph Theory*, 66(3):235–240, 2011.

[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.

[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.

[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.

[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.

[Tho88] Carsten Thomassen. On the presence of disjoint subgraphs of a specified type. *Journal of Graph Theory*, 12(1):101–111, 1988.

A Proof of Theorem 2.2

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 $\mathrm{tw}(G') = \Omega(k/\operatorname{poly}\log k)$.*

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 our input graph of treewidth at least k. We use Theorem A.1 to obtain a subgraph $G' \subseteq G$ of treewidth $k' = \Omega(k/\operatorname{poly}\log k)$ 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 2.2, 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 $C', S' \subseteq V(G)$, 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.1. *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 2.2. It now remains to prove Lemma A.1.

Proof of Lemma A.1. 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 subsetes $A_i' \subseteq A_i'', B_i' \subseteq B_i'', X_i' \subseteq X_i''$ 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 2.1, 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 $\mathcal{C}' = (C_1, C_3, \ldots, C_{2\lceil \ell/2 \rceil - 1})$, 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}'. Combining this new Path-of-Sets System with the clusters S_i and the sets $\{\mathcal{Q}_i\}$ of paths, for all odd indices $1 \leq i \leq \ell$, we obtain a hairy Path-of-Sets System of length $\ell'/2$ and width $\lceil w'/(3c_\Delta) \rceil$. $\square$