Layered Pathwidth and Graph Layouts

Vida Dujmović[1], Pat Morin[2], and Céline Yelle[1]

[1] School of Computer Science and EE, University of Ottawa, Ottawa, ON, Canada
{cyell045,vida.dujmovic}@uottawa.ca
[2] School of Computer Science, Carleton University, Ottawa, ON, Canada
morin@scs.carleton.ca

Abstract. We investigate a relationship between layered pathwidth and track number and layered pathwidth and stack number. It was known that graphs that have bounded layered pathwidth have bounded queue number and thus bounded track number (Bannister, Devanny, Dujmović, Eppstein, and Wood [GD 2015, Algorithmica 2018]). In the other direction, graphs with track number at most 2 are forests of caterpillars and have layered pathwidth at most 2. However, the existence of 4-track graphs that are expanders (Dujmović, Sidiropoulos and Wood [Chicago J. Theor. Comput. Sci. 2016]) implies that there are 4-track graphs whose layered pathwidth is $\Omega(n/\log n)$.
These results leave a gap for 3-track graphs. Indeed, Bannister et al. conjecture that 3-track graphs have bounded layered pathwidth. The main result of the current paper is to confirm this conjecture by showing that 3-track graphs have layered pathwidth at most 4. In a similar vein, the current paper also completes our knowledge about the relationship between stack number and layered pathwidth by showing that the stack number of a graph is at most 4 times its layered pathwidth.

1 Introduction

The treewidth and pathwidth of a graph are important tools in structural and algorithmic graph theory. Layered treewidth and layered H-partitions are recently developed tools that generalize treewidth. These tools played a critical role in recent breakthroughs on a number of problems including closing a 27 year old conjecture on queue layouts [10,9]. Motivated by the versatility and utility of layered treewidth, Bannister et al. [2,3] introduced layered pathwidth, which generalizes pathwidth in the same way that layered treewidth generalizes treewidth. The goal of this article is to fill the gaps in our knowledge about the relationship between layered pathwidth and the following well studied graph drawing layouts: queue-layouts, stack-layouts and track layouts.

A *tree decomposition* of a graph G is given by a tree T whose nodes index a collection of sets $B_1, \ldots, B_p \subseteq V(G)$ called *bags* such that (1) for each $v \in V(G)$, the set $T[v]$ of bags that contain v induce a non-empty (connected) subtree in T; and (2) for each edge $vw \in E(G)$, there is some bag that contains both v and w. If T is a path, the resulting decomposition is a *path decomposition.* The *width* of a tree (path) decomposition is the size of its largest bag. The *treewidth* (*pathwidth*) of G, denoted $\mathsf{tw}(G)$ ($\mathsf{pw}(G)$), is the minimum width of any tree (path) decomposition of G minus 1.

A *layering* of G is a mapping $\ell : V(G) \to \mathbb{Z}$ with the property that $vw \in E(G)$ implies $|\ell(u) - \ell(w)| \leq 1$. One can also think of a layering as a partition of G's vertices into sets indexed by integers, where $L_i = \{v \in V(G) : \ell(v) = i\}$ is called a *layer*. A *layered tree (path) decomposition* of G consists of a layering ℓ and a tree (path) decomposition with bags $B_1, \ldots, B_p$ of G. The *(layered) width* of a layered tree (path) decomposition is the maximum size of the intersection of a bag and a layer, i.e., $\mathsf{max}\{|L_i \cap B_j| : i \in \mathbb{Z},\ j \in \{1, \ldots, p\}\}$. The *layered treewidth (pathwidth)* of G, denoted $\mathsf{ltw}(G)$ ($\mathsf{lpw}(G)$) is the smallest (layered) width of any layered tree (path) decomposition of G.

Note that while layered pathwidth is at most pathwidth, pathwidth is not bounded by layered pathwidth. There are graphs—for example the $n \times n$ planar grid—that have unbounded pathwidth and bounded layered pathwidth. Thus upper bounds proved in terms of layered pathwidth are quantitatively stronger than those proved in terms of pathwidth. In addition while having pathwidth at most k is a minor-closed property, having layered pathwidth at most k is not. For example $2 \times n \times n$ grid graph has layered pathwidth at most 3 but it has K_n as a minor, and thus it has a minor of unbounded layered pathwidth. (Analogous statements hold for layered treewidth)

After introducing layered path-decompositions, Bannister et al. [2,3] set out to understand the relationship between track/queue/stack number and layered pathwidth (see the definitions in the footnote).[3] A summary of known and new results on these rich relationships are outlined in Table 1. The first two rows show (the older results) that track number and queue number are *tied*; each is bounded by some function of the other.

[3] A *t-track layout* of a graph G is a partition of $V(G)$ into t ordered independent sets $T_1, \ldots, T_t$ (with a total order $\prec_i$ for each T_i, $i \in \{1, \ldots, t\}$) with no X-crossings. Here an *X-crossing* is a pair of edges vw and xy such that, for some $i, j \in \{1, \ldots, t\}$, $v, x \in T_i$ with $v \prec_i x$ and $w, y \in T_j$ with $y \prec_j w$. The minimum number of tracks in any t-track layout of G is called the *track number* of G and is denoted as $\mathsf{tn}(G)$. A *t-track graph* is a graph that has a t-track layout.

A *stack (queue) layout* of a graph G consists of a total order σ of $V(G)$ and a partition of $E(G)$ into sets, called *stacks* (*queues*), such that no two edges in the same stack (queue) *cross*; that is, there are no edges vw and xy in a single stack with $v \prec_\sigma x \prec_\sigma w \prec_\sigma y$ (*nest*; there are no edges vw and xy in a single queue with $v \prec_\sigma x \prec_\sigma y \prec_\sigma w$.). The minimum number of stacks (queues) in a stack (queue) layout of G is the *stack number* (the *queue number*) of G and is denoted as $\mathsf{sn}(G)$ ($\mathsf{qn}(G)$). A stack layout is also called a *book embedding* and stack number is also called *book thickness* and *page number*. A *s-stack graph* (*q-queue graph*) is a graph that has a stack (queue) layout with at most s (q) queues.

Queue-Number versus Track-Number	
$\mathsf{qn}(G) \leq \mathsf{tn}(G) - 1$	[11, Theorem 2.6]
$\mathsf{tn}(G) \leq 2^{O(\mathsf{qn}(G)^2)}$	[12, Theorem 8]

Queue-Number versus Layered Pathwidth	
$\mathsf{qn}(G) \leq 3\,\mathsf{lpw}(G) - 1$	[11, Theorem 2.6][3, Lemma 9]
$\mathsf{qn}(G) = 1 \Rightarrow \mathsf{lpw}(G) \leq 2$	[15, Theorem 3.2][3, Corollary 7]
$\exists G : \mathsf{qn}(G) = 2,\ \mathsf{lpw}(G) = \Omega((n/\log n)$	[13, Theorem 1.4]

Track-Number versus Layered Pathwidth	
$\mathsf{tn}(G) \leq 3\,\mathsf{lpw}(G)$	[3, Lemma 9]
$\mathsf{tn}(G) = 1 \Rightarrow \mathsf{lpw}(G) = 1$	(G has no edges)
$\mathsf{tn}(G) \leq 2 \Rightarrow \mathsf{lpw}(G) \leq 2$	(G is a forest of caterpillars)
$\mathsf{tn}(G) \leq 3 \Rightarrow \mathsf{lpw}(G) \leq 4$	**Theorem 1**
$\exists G : \mathsf{tn}(G) = 4,\ \mathsf{lpw}(G) = \Omega((n/\log n)$	[13, Theorem 1.5]

Stack-Number versus Layered Pathwidth	
$\mathsf{sn}(G) \leq 1 \Rightarrow \mathsf{lpw}(G) \leq 2$	[3, Corollary 16]
$\exists G : \mathsf{sn}(G) = 2,\ \mathsf{lpw}(G) = \Omega(\log n)$	(G is a binary tree plus an apex vertex)
$\exists G : \mathsf{sn}(G) = 3,\ \mathsf{lpw}(G) = \Omega(n/\log n)$	[13, Theorem 1.5]
$\mathsf{sn}(G) \leq 4\,\mathsf{lpw}(G)$	**Theorem 2**

Table 1: Relationships between track number, queue number, stack number, and layered pathwidth.

The next group of rows relates queue number and layered pathwidth. Queue number is bounded by layered pathwidth. Graphs with queue number 1 are arched-level planar graphs and have layered pathwidth at most 2.[4] However, there are graphs with queue number 2 that are expanders; these graphs have pathwidth $\Omega(n)$ and diameter $O(\log n)$, so their layered pathwidth is $\Omega(n/\log n)$. Thus, layered pathwidth is not bounded by queue number.

The next group of rows relates track number and layered pathwidth. Track number is bounded by layered pathwidth. Layered pathwidth is bounded by track number when the track number is 1, or 2, but is not bounded by track number when the track number is 4 or more. The question of what happens for track number 3 is stated as an open problem by Bannister et al. [3], who solved the special case when G is bipartite and has track number 3. Our Theorem 1 solves this problem completely by showing that graphs with track number at most 3 have layered pathwidth at most 4.

Note that minor-closed classes that have bounded layered pathwidth have been characterized (as classes of graphs that exclude an apex tree[5] as a minor)

[4] Theorem 6 in [3] can easily be modified to prove that arched levelled planar graphs have layered pathwidth at most 2. That is achieved by placing the left most vertex from every level to every bag of the path decomposition.

[5] A graph G is an *apex tree* if it has a vertex v such that $G - v$ is a forest.

[7]. However, this result could not have been used to prove Theorem 1 since 3-track graphs are not closed under taking minors.[6]

Theorem 1. *Every graph G that has* $\mathsf{tn}(G) \leq 3$*, has* $\mathsf{lpw}(G) \leq 4$.

The final group of rows examines the relationship between stack number and layered pathwidth. Graphs of stack number at most 1 are exactly the outerplanar graphs, which have layered-pathwidth at most 2. On the other hand, there are graphs of stack number 2 that have unbounded layered pathwidth. Thus, in general, layered pathwidth is not bounded by stack number. Our second result, Theorem 2, shows that stack number is nevertheless bounded by layered pathwidth.

Theorem 2, for example, implies that unit-disk graphs with maximum clique size k have stack number $O(k)$ since they have been shown to have $O(k)$ layered pathwidth [2,3].

Theorem 2. *For every graph G,* $\mathsf{sn}(G) \leq 4\,\mathsf{lpw}(G)$.

We conclude this discussion by remarking that similar upper bounds for graph of bounded layered *treewidth* are not yet known, and present a challenging avenue for further study. For example, k-planar graphs are known to have layered treewidth $O(k)$ [8, Theorem 3.1]. Therefore, bounding stack number by a function of layered treewidth would imply that k-planar graphs have bounded stack-number. It is still unknown whether k-planar graphs have bounded stack-number except in the case $k = 1$ [6,1].

In the rest of the paper we sketch the proof of Theorem 1. The full proofs of Theorems 1 and 2 are included in the appendix.

2 Proof of Theorem 1

Let G be an edge-maximal n-vertex graph with $\mathsf{tn}(G) = 3$. Here, G is *edge-maximal* if adding an edge e creates a crossing. It is helpful to recall that G is a planar graph that has a straight-line crossing-free drawing with the vertices of T_1 placed on the positive x-axis, the vertices of T_2 placed on the positive y-axis and the vertices of T_3 placed on the ray $\{(a, a) : a < 0\}$. See Figure 1.

It will be easier to prove Theorem 1 for a weaker notion of layering. An *s-weak layering* of G is a mapping $\ell : V(G) \to \mathbb{Z}$ with the property that, for every $vw \in E(G)$, $|\ell(v) - \ell(w)| \leq s$. The sets $L_i = \{v \in V(G) : \ell(v) = i\}$ are called s-weak layers. The terms *s-weak layered path decomposition* and *s-weak layered pathwidth* of G, denoted $\mathsf{lpw}_s(G)$, are defined the same way as layered path decompositions and layered pathwidth, but with respect to s-weak layerings of G. The following result is easy (and well-known):

Lemma 1. *For any $s \in \mathbb{N}$,* $\mathsf{lpw}(G) \leq s \cdot \mathsf{lpw}_s(G)$.

[6] To see this, start with an $n \times n$ planar grid. Every planar grid has a 3-track layout. However, for large enough n, one can contract/delete edges on this grid graph such that the result is a series-parallel graph that does not have a 3-track layout (in particular the series-parallel graph from Theorem 18 in [3]).

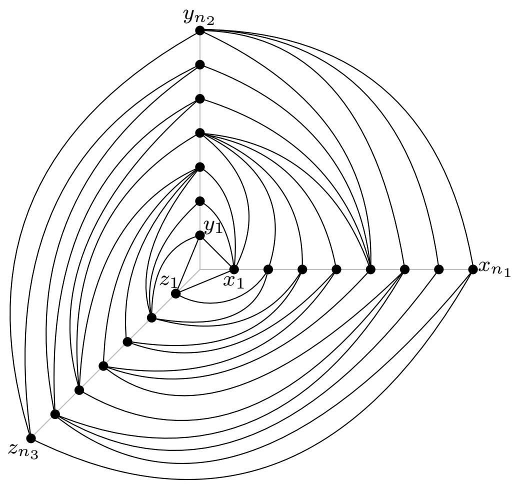

Fig. 1: The standard planar embedding of a 3-track graph.

Let T_1, T_2, T_3 be a 3-track layout of G with $T_1 = \{x_1, \ldots, x_{n_1}\}$, $T_2 = \{y_1, \ldots, y_{n_2}\}$, and $T_3 = \{z_1, \ldots, z_{n_3}\}$ and the total orders $\prec_1, \prec_2, \prec_3$ are implicit so that, for example $z_i \prec_3 z_j$ if and only if $i < j$. In terms of Figure 1, this means that x_1, y_1, z_1 form the triangular face containing the origin and $x_{n_1}, y_{n_2}, z_{n_3}$ form the cycle on the boundary of the outer face. From this point onward, all track indices are implicitly taken "modulo 3" so that for any integer i, T_i refers to the track $T_{i'}$ with index $i' = ((i-1) \bmod 3) + 1$.

Theorem 1 is a consequence of the following lemma.

Lemma 2. *The graph G described above has a 2-weak layered path decomposition, $B_1, \ldots, B_p$, with layering ℓ of (layered) pathwidth 2 in which*

1. *for each $i \in \{1, 2, 3\}$ and each $v \in T_i$, $\ell(v) \equiv i \pmod 3$;*
2. *$B_1 = \{x_1, y_1, z_1\}$;*
3. *$\ell(x_1) = 1$, $\ell(y_1) = 2$, and $\ell(z_1) = 3$;*
4. *$B_p = \{x_{n_1}, y_{n_2}, z_{n_3}\}$; and*
5. *$x_{n_1}, y_{n_2}, z_{n_3}$ appear in 3 distinct consecutive layers.*

Before proving Lemma 2, we show how it implies Theorem 1. Since layered pathwidth is monotone with respect to the addition of edges, it is safe to assume (as Lemma 2 does) that G is edge-maximal. By Lemma 2, therefore G has $\mathsf{lpw}_2(G) \leq 2$ and therefore, by Lemma 1, $\mathsf{lpw}(G) \leq 4$.

Sketch of proof of Lemma 2. The proof is by induction on the number $|V(G)|$ of vertices. If $|V(G)| \leq 4$, then the result is trivial.

Suppose that G has a cut set $C = \{x_i, y_j, z_k\}$ having exactly one vertex in each track. Since G is edge-maximal, x_i, y_j, z_k form a cycle in G. Now, the subgraph G_1 of G induced by $\{x_1, \ldots, x_i, y_1, \ldots, y_j, z_1, \ldots, z_k\}$ is an edge-maximal graph with $\mathsf{tn}(G_1) = 3$. Similarly, the subgraph G_2 of G induced by $\{x_i, \ldots, x_{n_1}, y_j, \ldots, y_{n_2}, z_k, \ldots, z_{n_3}\}$ is edge-maximal with $\mathsf{tn}(G_2) = 3$. We can therefore apply the lemma inductively on each of G_1 and G_2. By appropriately relabelling x, y, and z in G_2, Conditions 2 and 4 of the lemma ensure that the desired path decomposition of G can be obtained by concatenating the path decompositions of G_1 and G_2. Conditions 3 and 5 ensure that a layering of G can be obtained by combining the layering of G_1 with (a shifted version) the layering of G_2.

Thus, we can focus on case where G has no separating triangle. In this case, we greedily construct a path $P = v_1, ..., v_r$. The first vertex of P will be one of x_1, y_1, z_1 and the last three vertices will be $x_{n_1}, y_{n_2}, z_{n_3}$. For any $v_k \in P$, we select the next vertex v_{k+1} by choosing the neighbouring vertex of v_k on track T_{k+1} with the highest index. See Figure 2 for an example of the constructed path P.

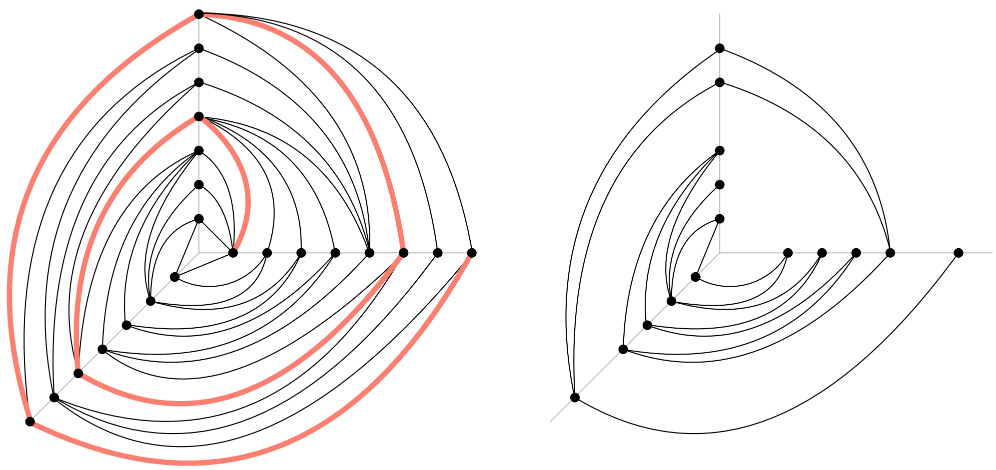

Fig. 2: The graph $G - P$ is a levelled planar graph.

The 2-weak layering of G is obtained as follows: For each vertex v_i on P, we will set $\ell(v_i) = i$ For each $t \in \{1, 2, 3\}$, any vertex $v \in T_t$ that is not in P is assigned to the same layer as v's immediate successor in $P \cap T_t$. That this is a 2-weak layering follows from straightforward analysis. (The only edges that span 2 layers are those incident to a vertex in $v \in T_{i-1} \setminus V(P)$ and a vertex $v_i \in T_i \cap V(P)$, where v is further from the origin than v_{i-1}.)

Next, we consider the graph $G - P$ obtained by removing the vertices of P from G. Clearly $\mathsf{tn}(G - P) \leq 3$ and, since every edge of $G - P$ joins a pair of vertices in two consecutive layers, every cycle in $G - P$ has even length, so $G - P$ is bipartite. By a result of Bannister et al. [3, Proof of Theorem 5], $G - P$ has a layered path decomposition $B_1, \ldots, B_p$ of width 1 using the layering ℓ defined above. If we add the vertices of P to every bag of this decomposition we obtain a width-2 2-weak layered path decomposition of G. Finally, to satisfy Conditions 2 and 4 of the lemma, we prepend a bag $B_0 = \{x_1, y_1, z_1\}$ and append a bag $B_{p+1} = \{x_{n_1}, y_{n_2}, z_{n_3}\}$. $\square$

References

1. Md. Jawaherul Alam, Franz J. Brandenburg, and Stephen G. Kobourov. On the book thickness of 1-planar graphs. *CoRR*, abs/1510.05891, 2015.
2. Michael J. Bannister, William E. Devanny, Vida Dujmović, David Eppstein, and David R. Wood. Track layout is hard. In *Graph Drawing and Network Visualization - 24th Int. Symposium, GD 2016*, volume 9801 of *LNCS*, pages 499–510, 2016.
3. Michael J. Bannister, William E. Devanny, Vida Dujmović, David Eppstein, and David R. Wood. Track layouts, layered path decompositions, and leveled planarity. *Algorithmica*, pages 1–23, July 2018.
4. Oliver Bastert and Christian Matuszewski. Layered drawings of digraphs. *Drawing Graphs, Methods and Models*, 2025:87–120, 2001.
5. Giuseppe Di Battista, Peter Eades, Roberto Tamassia, and Ioannis G. Tollis. Layered drawings of digraphs. *Graph Drawing: Algorithms for the Visualization of Graphs*, pages 265–302, 1999.
6. Michael A. Bekos, Till Bruckdorfer, Michael Kaufmann, and Chrysanthi N. Raftopoulou. The book thickness of 1-planar graphs is constant. *Algorithmica*, 79(2):444–465, 2017.
7. Vida Dujmović, David Eppstein, Gwenaël Joret, Pat Morin, and David R. Wood. Minor-closed graph classes with bounded layered pathwidth. *CoRR*, abs/1810.08314, 2018.
8. Vida Dujmović, David Eppstein, and David R. Wood. Structure of graphs with locally restricted crossings. *SIAM J. Discrete Math.*, 31(2):805–824, 2017.
9. Vida Dujmovic, Louis Esperet, Gwenaël Joret, Bartosz Walczak, and David R. Wood. Planar graphs have bounded nonrepetitive chromatic number. *CoRR*, abs/1904.05269, 2019.
10. Vida Dujmovic, Gwenaël Joret, Piotr Micek, Pat Morin, Torsten Ueckerdt, and David R. Wood. Planar graphs have bounded queue-number. *CoRR*, abs/1904.04791, 2019.
11. Vida Dujmovic, Pat Morin, and David R. Wood. Layout of graphs with bounded tree-width. *SIAM J. Comput.*, 34(3):553–579, 2005.
12. Vida Dujmovic, Attila Pór, and David R. Wood. Track layouts of graphs. *Discrete Mathematics & Theoretical Computer Science*, 6(2):497–522, 2004.
13. Vida Dujmović, Anastasios Sidiropoulos, and David R. Wood. Layouts of expander graphs. *Chicago J. Theor. Comput. Sci.*, 2016, 2016.
14. Patrick Healy and Nikola S. Nikolov. Hierarchical drawing algorithms. *Handbook on Graph Drawing and Visualization*, pages 409–453, 2013.
15. Lenwood S. Heath and Arnold L. Rosenberg. Laying out graphs using queues. *SIAM J. Comput.*, 21(5):927–958, 1992.
16. K. Sugiyama, S. Tagawa, and M. Toda. Methods for visual understanding of hierarchical system structures. *IEEE Trans. Systems Man Cybernet.*, 11(2):109–125, 1981.

A Proof of Lemma 2

The following observation follows from the fact that G is edge-maximal.

Observation 1. *For any two vertices of G on distinct tracks, say x_i and y_j, at least one of the following conditions is satisfied (see Figure 3):*

1. $x_i y_j \in E(G)$; *or*
2. *there exists* $x_{i'} y_{j'} \in E(G)$ *with* $i' > i$ *and* $j' < j$; *or*
3. *there exists* $x_i y_{j'} \in E(G)$ *with* $j' > j$.

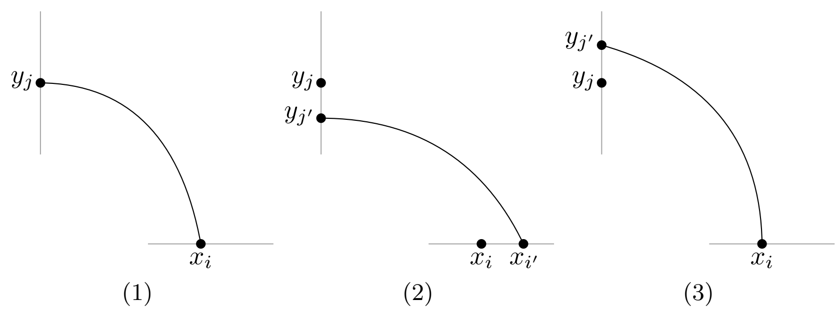

Fig. 3: The three cases in Observation 1.

Lemma 2. *The graph G described above has a 2-weak layered path decomposition, $B_1, \dots, B_p$, with layering ℓ of (layered) pathwidth 2 in which*

1. *for each $i \in \{1, 2, 3\}$ and each $v \in T_i$, $\ell(v) \equiv i \pmod 3$;*
2. *$B_1 = \{x_1, y_1, z_1\}$;*
3. *$\ell(x_1) = 1$, $\ell(y_1) = 2$, and $\ell(z_1) = 3$;*
4. *$B_p = \{x_{n_1}, y_{n_2}, z_{n_3}\}$; and*
5. *$x_{n_1}, y_{n_2}, z_{n_3}$ appear in 3 distinct consecutive layers.*

Proof of Lemma 2. The proof is by induction on the number of vertices of G. If $|V(G)| \leq 4$, then the result is trivial. Next, suppose that G has a cut set $C = \{x_i, y_j, z_k\}$ having exactly one vertex in each track. Since G is edge-maximal, x_i, y_j, z_k form a cycle in G. Now, the subgraph G_1 of G induced by $\{x_1, \dots, x_i, y_1, \dots, y_j, z_1, \dots, z_k\}$ is an edge-maximal graph with $\mathsf{tn}(G_1) = 3$ and we can inductively apply Lemma 2 to find a width-2 2-weak layered path decomposition of G_1 in which x_i, y_j, z_k are in the last bag and are assigned to three consecutive distinct layers $r+1$, $r+2$, and $r+3$. Note that there are three possible assignments of x_i, y_j, z_k to these three layers depending on the value of $r \bmod 3$. Suppose, without loss of generality, that $\ell(y_j) = r+1$ (so $\ell(z_k) = r+2$ and $\ell(x_i) = r+3$.)

Next, consider the graph G_2 induced by $\{x_i, \dots, x_{n_1}, y_j, \dots, y_{n_2}, z_k, \dots, z_{n_3}\}$. We apply Lemma 2 inductively on G_2 relabelling tracks to ensure that in the resulting layered decomposition $\ell(y_j) = 1$, $\ell(z_k) = 2$ and $\ell(x_i) = 3$. We can now obtain a width-2 2-weak layered path decomposition of G by joining the two decompositions. In particular, concatenating the sequence of bags for G_1 with the sequence of bags for G_2 gives a path decomposition of G and adding r to the indices of all layers in the layering of G_2 gives a 2-weak layering of G.

Thus, all that remains is to study the case where G contains no cut set having exactly one vertex on each track. We claim that, in this case, G contains the edge $x_1 z_2$ or it contains the edge $z_1 x_2$. Since G is edge-maximal, this is obvious unless $n_1 = n_3 = 1$ so that neither z_2 nor x_2 exist. However, since $|V(G)| \geq 5$, this would imply that x_1, z_1, y_2 is a cut set with one vertex on each track, since its removal separates all y_i, with $i \neq 2$.

We will construct a path $P = v_1, \ldots, v_r$, an example of which is illustrated in Figure 4. The first vertex of P will be one of x_1, y_1, z_1 and the final three vertices are x_{n_1}, y_{n_2}, and z_{n_3}, though not necessarily in that order. The path P will *spiral* in the sense that $v_i \in T_i$ for all $i \in \{1, \ldots, r\}$. Observe that this spiralling implies that the subsequence of vertices of P on any track T_i is increasing (getting further from the origin).

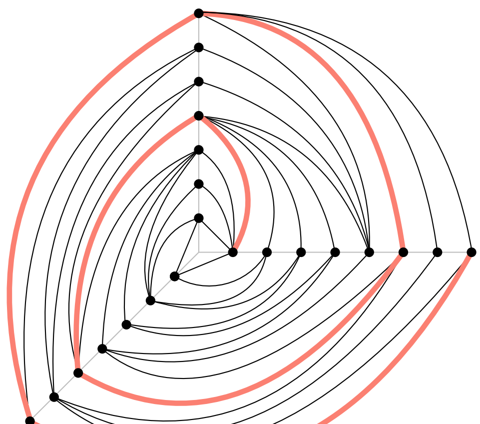

Fig. 4: The path P used in the proof of Lemma 2.

P is constructed greedily: if P has reached vertex v_k, it continues to the neighbouring vertex v_{k+1} of v_k with the highest index on track T_{k+1} that is not yet in P. We will call this vertex v_{k+1} the *furthest neighbouring vertex* of v_k. To see why this is always possible, recall that P already contains an edge v_{k-3}, v_{k-2}. Now, without loss of generality we can apply Observation 1 with $x_i = v_k$ and $y_j = v_{k-2}$, so there are three cases (see Figure 5):

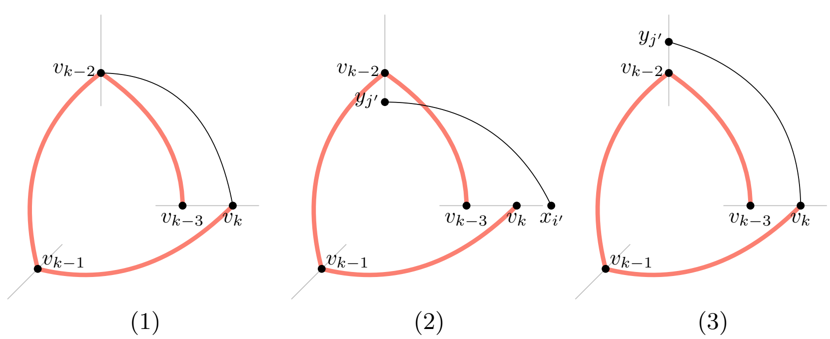

Fig. 5: The path P can always be extended.

1. $v_k v_{k-2} \in E(G)$. In this case v_{k-2}, v_{k-1}, and v_k form a cycle in G. Then either $\{v_{k-2}, v_{k-1}, v_k\} = \{x_{n_1}, y_{n_2}, z_{n_3}\}$ or $\{v_{k-2}, v_{k-1}, v_k\}$ is a cut set with exactly one vertex in each track. In the former case, the path P is complete. The latter case is excluded by the assumption that G contains no such cut sets.
2. there exists an edge $x_{i'} y_{j'} \in E(G)$ with $i' > i$ (i.e. $i' > k$) and $j' < j$ (i.e. $j' < k-2$). This case is not possible, since this edge would cross $v_{k-3} v_{k-2}$.
3. there exists an edge $v_k y_{j'} \in E(G)$ with $j' > j$ (i.e. $j' > k-2$). In this case, P is extended by adding $v_{k+1} = y_{j'}$.

Thus we have constructed the furthest vertex path $P = v_1, \dots, v_r$ whose first vertex is one of x_1, y_1, z_1 and whose last three vertices are x_{n_1}, y_{n_2} and z_{n_3} (not necessarily in order). We assign layers to the vertices of P as follows: For each vertex v_i on P, we set $\ell(v_i) = i$. Note that this satisfies Conditions 3 and 5 of the lemma and also satisfies Condition 1 for the vertices of P. For each $t \in \{1, 2, 3\}$, any vertex $v \in T_t$ that is not in P is assigned to the same layer as v's immediate successor in $P \cap T_t$. This assignment satisfies Condition 1 for vertices not in P. Finally, we will prove that this gives a 2-weak layering of G. In other words, in the worst case, a vertex v with $\ell(v) = i$ can only share an edge with vertex u where $i - 2 \le \ell(u) \le i + 2$.

Any edge between v and w where neither v nor w is in P will only span one layer. Any edge between any two vertices v_i and v_j where $v_i, v_j \in P$, will span only one layer if $j = i \pm 1$. This would mean that $v_i v_j$ is an edge in the graph G and that this edge was used to construct our furthest vertex path P. If $j \ne i \pm 1$, then there are two cases:

1. $j = i \pm 2$ This edge is possible and would create a cut set. This edges will only span 2 layers since $\ell(v_i) = i$ and $\ell(v_j) = i + 2$.
2. $j = i \pm 4$ This edge cannot exist since it would contradict our greedy path constructing algorithm. If the edge $v_i v_{i+4}$ (or the edge $v_{i-4} v_i$) existed then the edge $v_i v_{i+1}$ ($v_{i-4} v_{i-3}$) would not have been added to P.

Any edge between v and w where $v \in P$ and $w \notin P$ will have 7 cases (see Figure 6). Without loss of generality, assume the spiral is travelling from T_1 to T_2 to T_3. Let x_i be a vertex on the constructed path P. First, we look at the possible cases for an edge between x_i with $\ell(x_i) = m$ and y_j where $y_j \notin P$.

1. $\ell(y_j) = m + 3n$ where $n \ge 1$. This edge cannot exist since it would contradict our greedy path constructing algorithm.
2. $\ell(y_j) = m + 1$. This edge will only span one layer.
3. $\ell(y_j) = m + 1 - 3n$ where $n \ge 1$. This edge cannot exist, since it would create a crossing with the edge $v_{m-3} v_{m-2}$.

Second, we look at the possible cases for an edge between x_i with $\ell(x_i) = m$ and z_k where $z_k \notin P$.

4. $\ell(z_k) = m + 2 + 3n$ where $n \ge 1$. This edge cannot exist, since it would create a crossing with the edge $v_{m+2} v_{m+3}$.

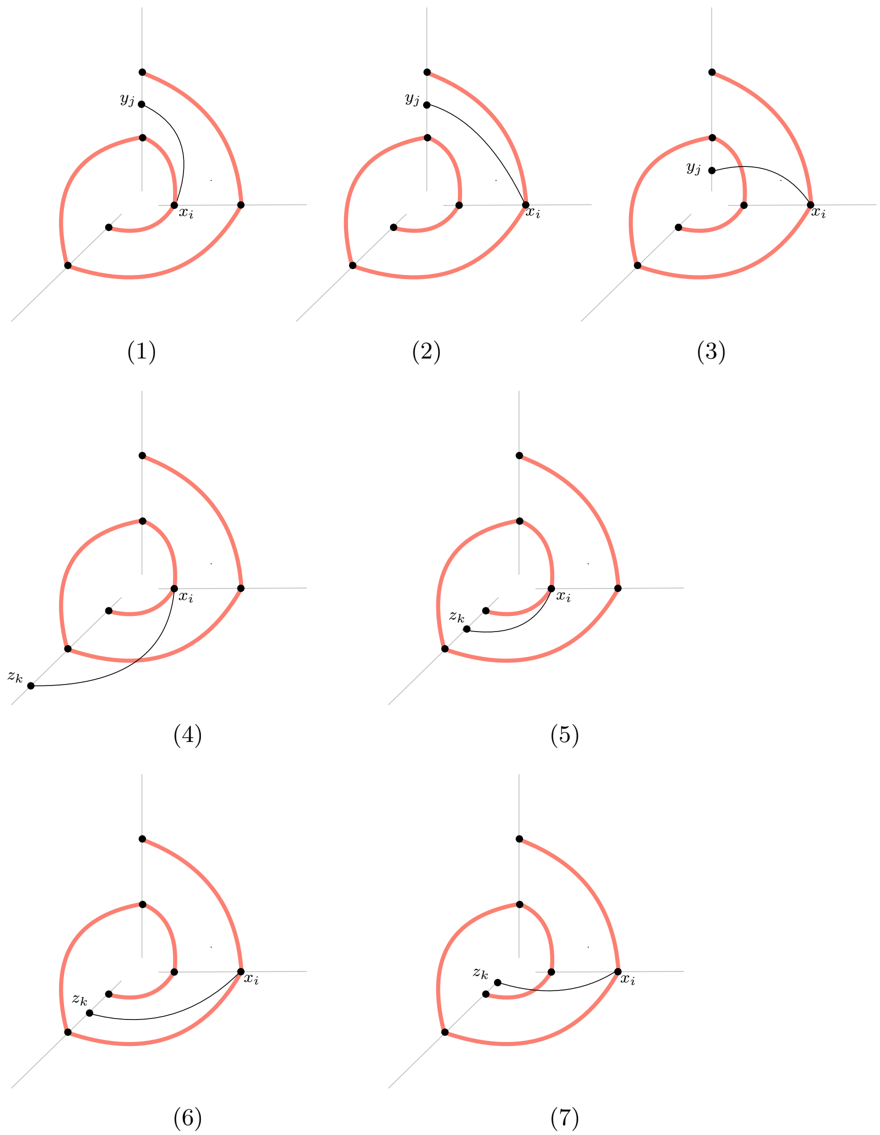

Fig. 6: The edge between a vertex x_i and a vertex y_j or z_k cannot span more than 2 layers

5. $\ell(z_k) = m + 2$. This edge spans exactly two layers.
6. $\ell(z_k) = m - 1$. This edge will only span one layer.
7. $\ell(z_k) = m - 1 - 3n$ where $n \geq 1$. This edge cannot exist, since it would create a crossing with the edge $v_{m-4}v_{m-3}$.

Next, we will need a notion of levelled planar graphs. The class of levelled planar graphs was introduced in 1992 by Heath and Rosenberg [15] in their study of queue layouts of graphs. A levelled planar drawing of a graph is a straight-line crossing-free drawing in the plane, such that the vertices are placed on a sequence of parallel lines (called levels), where each edge joins vertices in two consecutive levels. A graph is levelled planar if it has a levelled planar drawing. (This is a well studied model for planar graph drawing, so called Sugiyama-style [16,4,14,5].)

Now, consider the graph $G - P$ obtained by removing the vertices of P from G (see Figure 2). We claim that this graph is a levelled planar graph in which the levels of the vertices are given by the layering ℓ defined above. Refer to Figure 7. One way to see this is to imagine G as being drawn with its vertices on the three vertical edges of a hollow triangular prism so that x_1, y_1, z_1 are the vertices of one triangular face and $x_{n_1}, y_{n_2}, z_{n_3}$ are the vertices of the other triangular face. Now, if we remove the triangular faces of this prism, cut it along the embedding of P, and unfold the resulting surface so that it lies in the plane, then we obtain a drawing of $G - P$ in which the vertices lie on a set of parallel lines and in which the edges only join vertices on two consecutive lines. This gives the desired levelled planar drawing of $G - P$.

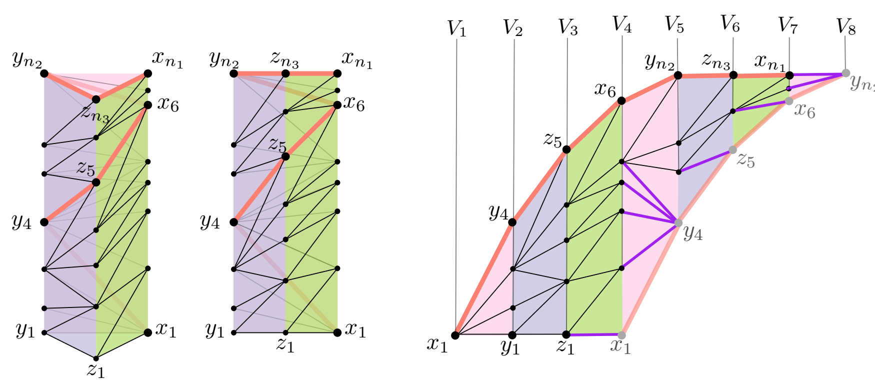

Fig. 7: Cutting a prism along P to obtain a levelled planar drawing of $G - P$. Edges that span 2 layers are drawn in purple.

By a result of Bannister et al. [3, Proof of Theorem 5], $G - P$ has a layered path decomposition $B_1, \ldots, B_p$ of width 1 using the layering ℓ defined above. If we add the vertices of P to every bag of this decomposition we obtain a width-2 2-weak layered path decomposition of G. Finally, to satisfy Conditions 2 and 4 of the lemma, we prepend a bag $B_0 = \{x_1, y_1, z_1\}$ and append a bag $B_{p+1} = \{x_{n_1}, y_{n_2}, z_{n_3}\}$. □

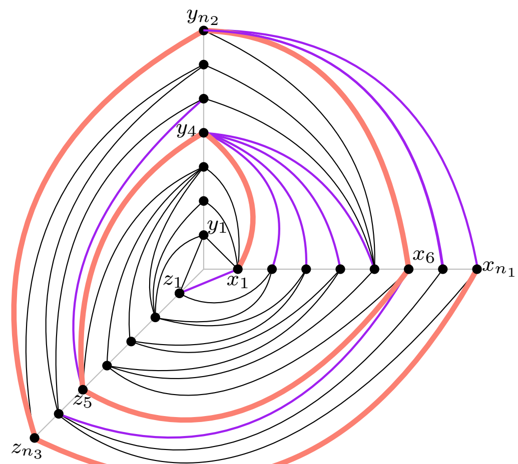

Fig. 8: The graph G with the path P and edges that span 2 layers drawn in purple

B Proof of Theorem 2

Let $B = B_1, \dots, B_p$ be a path decomposition of G and $\ell : V(G) \to \mathbb{Z}$ be a layering that obtains so that B has layered width $\mathsf{lpw}(G)$ with respect to the layering ℓ.

We may assume that B is *left-normal* in the sense that, for every distinct pair $v, w \in V(G)$, $\min\{i \in \mathbb{Z} : v \in B_i\} \neq \min\{i \in \mathbb{Z} : w \in B_i\}$. It is straightforward to verify that any path decomposition can be made left-normal without increasing the layered width of the decomposition. We use the notation $v \prec_B w$ if $\min\{i \in \mathbb{Z} : v \in B_i\} < \min\{i \in \mathbb{Z} : w \in B_i\}$. Since B is left-normal, $\prec_B$ is a total order.

For any edge vw with $v \prec_B w$ we call v the *left endpoint* of the edge and w the *right endpoint*. We use the convention of writing any edge with endpoints v and w as vw where v is the left endpoint and w is the right endpoint. With these definitions and conventions in hand, we are ready to proceed.

Proof of Theorem 2. Consider a left-normal path decomposition $B = B_1, \dots, B_p$ of G and a layering $\ell : V(G) \to \mathbb{Z}$ such that B has layered with at most k with respect to ℓ. Thus, our goal is to construct a stack layout of G using $4k$ stacks.

We first construct total total ordering $\prec_\sigma$ on the vertices of G, as follows:

- If $v \in L_i$ and $w \in L_j$ with $i < j$, then $v \prec_\sigma w$.
- If $v, w \in L_i$ with $v \prec_\sigma w$ then
  - $v \prec_\sigma w$ if i is even; or
  - $w \prec_\sigma v$ if i is odd.

Next we define a colouring $\varphi : E(G) \to \{0,1\} \times \{0,1\} \times \{1, \dots, k\}$ that determines the partition of the $E(G)$ into stacks. We begin with a (greedy) vertex k-colouring $\varphi : V \to \{1, \dots, k\}$ so that, for any $i, j \in \mathbb{N}$, no two vertices in $B_i \cap L_j$ are assigned the same colour. This is easily done since the path decomposition $\langle B_i \cap L_j : i \in \mathbb{Z} \rangle$ has bags of size at most k.

We say that the edge vw has *positive slope* if $\ell(v) = \ell(w) + 1$ and has *non-positive slope* otherwise. We colour the edge vw with the colour $\varphi(vw) = (a, b, c)$ where $a = \ell(v) \bmod 2$, b is a bit indicating whether vw has positive or non-positive slope, and c is the colour $\varphi(v)$ of the left endpoint v. This clearly uses only $2 \times 2 \times k = 4k$ colours so all that remains is to show that σ and the partition $P = \{\{vw \in E(G) : \varphi(vw) = (a, b, c)\} : (a, b, c) \in \{0, 1\} \times \{0, 1\} \times \{1, \dots, k\}\}$ is indeed a stack layout.

Consider any two distinct edges $vw, xy \in E(G)$ (whose left-endpoints are v and x, respectively). First observe that, if $\ell(v) \equiv \ell(x) \pmod 2$ then either $\ell(v) = \ell(x)$ or $\ell(v) - \ell(x) \geq 2$. In the latter case, the only way in which vw and xy can cross with respect to $\prec_\sigma$ is if $\ell(v) + b = \ell(y) = \ell(w) = \ell(x) - b$ for some $b \in \{-1, 1\}$. However, in this case, vw has positive slope and xy has non-positive slope, or vice-versa, so $\varphi(vw)$ and $\varphi(xy)$ differ in their second component.

Therefore, we only need to consider pairs of edges xy and vw where $\ell(v) = \ell(x) = i$. We assume, without loss of generality that i is even and that $v \prec_\sigma x$. With these assumptions, there are only three cases in which vw and xy can cross:

1. $v \prec_\sigma x \prec_\sigma w \prec_\sigma y$. Since $\ell(v) = \ell(x) = i$ is even and $v \prec_\sigma x$, we have $v \prec_B x$ and $\ell(w) \geq i$. If $\ell(w) = i$, then $v \prec_B x \prec_B w$, so v and x both appear in some bag B_j and $\varphi(v) \neq \varphi(x)$, so $\varphi(vw)$ and $\varphi(xy)$ differ in their third component. If $\ell(w) = i + 1$, then $w \prec_\sigma y$ implies that $\ell(y) \geq \ell(w)$, which implies $\ell(y) = \ell(w) = i + 1$, so $y \prec_B w$. We now have $v \prec_B x \prec_B y \prec_B w$ so v and x appear in a common bag B_j and $\varphi(vw)$ and $\varphi(xy)$ differ in their third component.
2. $v \prec_\sigma y \prec_\sigma w \prec_\sigma x$. Since $v \prec_\sigma y$, $\ell(y) \geq \ell(v) = i$. Similarly, since $y \prec_\sigma x$, $\ell(y) \leq \ell(x) = i$. Therefore, $\ell(y) = i$, so $y \prec_B x$. This is not possible since, by definition, x is the left endpoint of xy.
3. $y \prec_\sigma v \prec_\sigma x \prec_\sigma w$. Since $y \prec_\sigma x$, x is the left endpoint of xy, and $\ell(x) = i$ is even, we have $\ell(y) = \ell(x) - 1$, so xy has positive slope.
   Since $v \prec_\sigma w$ and $\ell(v) = i$ is even, we have $\ell(v) \leq \ell(w)$, so vw has non-positive slope. Therefore $\varphi(vw)$ and $\varphi(xy)$ differ in their second component.

Therefore, for any pair of edges $vw, xy \in E(G)$ that cross, $\varphi(vw) \neq \varphi(xy)$, so the partition P is a partition of $V(G)$ into $4k$ stacks with respect to $\prec_\sigma$, as required. $\square$