Twin-Width of Planar Graphs Is at Most 8, and at Most 6 When Bipartite Planar

Petr Hliněný
Masaryk University, Brno, Czech republic

Jan Jedelský
Masaryk University, Brno, Czech republic

Abstract

Twin-width is a structural width parameter introduced by Bonnet, Kim, Thomassé and Watrigant [FOCS 2020]. Very briefly, its essence is a gradual reduction (a contraction sequence) of the given graph down to a single vertex while maintaining limited difference of neighbourhoods of the vertices, and it can be seen as widely generalizing several other traditional structural parameters. Having such a sequence at hand allows us to solve many otherwise hard problems efficiently. Graph classes of bounded twin-width, in which appropriate contraction sequences are efficiently constructible, are thus of interest in combinatorics and in computer science. However, we currently do not know in general how to obtain a witnessing contraction sequence of low width efficiently, and published upper bounds on the twin-width in non-trivial cases are often "astronomically large".

We focus on planar graphs, which are known to have bounded twin-width (already since the introduction of twin-width), but the first explicit "non-astronomical" upper bounds on the twin-width of planar graphs appeared just a year ago; namely the bound of at most 183 by Jacob and Pilipczuk [arXiv, January 2022], and 583 by Bonnet, Kwon and Wood [arXiv, February 2022]. Subsequent arXiv manuscripts in 2022 improved the bound down to 37 (Bekos et al.), 11 and 9 (both by Hliněný). We further elaborate on the approach used in the latter manuscripts, proving that the twin-width of every planar graph is at most 8, and construct a witnessing contraction sequence in linear time. Note that the currently best lower-bound planar example is of twin-width 7, by Král' and Lamaison [arXiv, September 2022]. We also prove that the twin-width of every bipartite planar graph is at most 6, and again construct a witnessing contraction sequence in linear time.

2012 ACM Subject Classification Mathematics of computing → Graph theory

Keywords and phrases twin-width, planar graph

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.75

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2210.08620` [17]

Funding Jan Jedelský was partially supported by Masaryk University project MUNI/I/1677/2018.

1 Introduction

Twin-width is a relatively new structural width measure of graphs and relational structures introduced in 2020 by Bonnet, Kim, Thomassé and Watrigant [10]. Informally, twin-width of a graph measures how diverse the neighbourhoods of the graph vertices are. E.g., *cographs* – the graphs which can be built from singleton vertices by repeated operations of a disjoint union and taking the complement, have the lowest possible value of twin-width, 0, which means that the graph can be brought down to a single vertex by successively identifying twin vertices. (Two vertices x and y are called *twins* in a graph G if they have the same neighbours in $V(G) \setminus \{x, y\}$.) Hence the name, *twin-width*, for the parameter.

Importance of this new concept is clearly witnessed by numerous recent papers on the topic, such as the follow-up series [5–9, 11] and more related research papers represented by, e.g., [1, 2, 4, 13, 15, 20].

Twin-width definition. In general, the concept of twin-width can be considered over arbitrary binary relational structures of a finite signature, but here we will define it and deal with it for only finite *simple graphs*, i.e., graphs without loops and multiple edges. A *trigraph* is a simple graph G in which some edges are marked as *red*, and with respect to the red edges only, we naturally speak about *red neighbours* and *red degree* in G. However, when speaking about edges, neighbours and/or subgraphs without further specification, we count both ordinary and red edges together as one edge set denoted by $E(G)$. The edges of G which are not red are sometimes called (and depicted) black for distinction. For a pair of (possibly not adjacent) vertices $x_1, x_2 \in V(G)$, we define a *contraction* of the pair x_1, x_2 as the operation creating a trigraph G' which is the same as G except that x_1, x_2 are replaced with a new vertex x_0 (said to *stem from* x_1, x_2) such that:

- the (full) neighbourhood of x_0 in G' (i.e., including the red neighbours), denoted by $N_{G'}(x_0)$, equals the union of the neighbourhoods $N_G(x_1)$ of x_1 and $N_G(x_2)$ of x_2 in G except x_1, x_2 themselves, that is, $N_{G'}(x_0) = (N_G(x_1) \cup N_G(x_2)) \setminus \{x_1, x_2\}$, and
- the red neighbours of x_0, denoted here by $N^r_{G'}(x_0)$, inherit all red neighbours of x_1 and of x_2 and add those in $N_G(x_1)\Delta N_G(x_2)$, that is, $N^r_{G'}(x_0) = \big(N^r_G(x_1) \cup N^r_G(x_2) \cup (N_G(x_1)\Delta N_G(x_2))\big) \setminus \{x_1, x_2\}$, where Δ denotes the symmetric set difference.

A *contraction sequence* of a trigraph G is a sequence of successive contractions turning G into a single vertex, and its *width* d is the maximum red degree of any vertex in any trigraph of the sequence. We also then say that it is a d-contraction sequence of G. The *twin-width* of a trigraph G is the minimum width over all possible contraction sequences of G. In other words, a graph has twin-width at most d, if and only if it admits a d-contraction sequence.

To define the twin-width of an ordinary (simple) graph G, we consider G as a trigraph with no red edges.

Algorithmic aspects. Twin-width, as a structural width parameter, has naturally many algorithmic applications in the FPT area. Among the most important ones we mention that the first order (FO) model checking problem – that is, deciding whether a fixed first-order sentence holds in an input graph – can be solved in linear FPT-time [11]. This and other algorithmic applications assume that a contraction sequence of bounded width is given alongside with the input graph. Deciding the exact value of twin-width (in particular, twin-width 4) is in general NP-hard [4], but for many natural graph classes we know that they are of bounded twin-width. However, published upper bounds on the twin-width in non-trivial cases are often non-explicit or "astronomically large", and it is not usual that we could, alongside such a bound, compute a contraction sequence of provably "reasonably small" width efficiently and practically. We pay attention to this particular aspect; and we will accompany our fine mathematical upper bounds on the twin-width with rather simple linear-time algorithms for computing contraction sequences of the claimed widths.

Twin-width of planar graphs. The fact that the class of planar graphs is of bounded twin-width was mentioned already in the pioneering paper [10], but without giving any explicit upper bound on the twin-width. The first explicit (numeric) upper bounds on the twin-width of planar graphs have been published only quite recently; chronologically on arXiv, the bound of 183 by Jacob and Pilipczuk [18], of 583 by Bonnet, Kwon and Wood [12] (this paper more generally bounds the twin-width of k-planar graphs by asymptotic $2^{\mathcal{O}(k)}$), and of 37 by Bekos, Da Lozzo, Hliněný, and Kaufmann [3] (this paper more generally bounds the twin-width of so-called h-framed graphs by $\mathcal{O}(h)$).

It is worth to mention that all three papers [3, 12, 18], more or less explicitly, use the product structure machinery of planar graphs (cf. [14]). We have then developed an alternative decomposition-based approach, leading to a single-digit upper bound of 9 for all planar graphs in [16], followed by an upper bound of 6 on the twin-width of *bipartite* planar graphs thereafter. However, the approach of [16] seems to be stuck right at 9, and new ideas were needed to obtain further improvements, even by 1.

Here we give the following strengthened upper bound, which uses an improved approach over previous [16] and also simplifies some cumbersome technical details of the former:

▶ **Theorem 1.** *The twin-width of any simple planar graph is at most* 8*, and a corresponding contraction sequence can be found in linear time.*

It is worth to note that, recently, Král' and Lamaison [19] have found a construction (with a proof) of a planar graph with twin-width 7. Hence, the lower bound is by just one off our upper bound, but the right maximum value (7 or 8?) is still an open question.

In an addition, we also prove an upper bound for bipartite planar graphs, which follows from an adaptation of our new techniques specially to the bipartite case. While the bipartite-case bound stays the same as in previous [16], its proof is significantly simpler now.

▶ **Theorem 2.** *The twin-width of any simple* bipartite *planar graph is at most* 6*, and a corresponding contraction sequence can be found in linear time.*

Due to space restrictions, proofs of the *-marked statements are left for the full paper [17].

2 Notation and Tools

We start with a few technical definitions and claims needed for the proofs.

BFS layering and contractions. Let G be a *connected* graph and $r \in V(G)$ a fixed vertex. The *BFS layering of G determined by r* is the vertex partition $\mathcal{L} = (L_0 = \{r\}, L_1, L_2, \dots)$ of G such that L_i contains all vertices of G at distance exactly i from r. A path $P \subseteq G$ is *r-geodesic* if P is a subpath of some shortest path from r to any vertex of G (in particular, P intersects every layer of $\mathcal{L}$ in at most one vertex). Let T be a *BFS tree* of G rooted at the vertex r as above (that is, for every vertex $v \in V(G)$, the distance from v to r is the same in G as in T). A path $P \subseteq G$ is *T-vertical*, or shortly *vertical* with respect to implicit T, if P is a subpath of some root-to-leaf path of T. Notice that a T-vertical path is r-geodesic, but the converse may not be true. Analogously, an edge $e \in E(G)$ is *horizontal* (with respect to implicit $\mathcal{L}$) if both ends of G are in the same $\mathcal{L}$-layer.

Observe the following trivial claim (cf. also Claim 5):

▷ Claim 3. For every edge $\{v, w\}$ of G with $v \in L_i$ and $w \in L_j$, we have $|i - j| \leq 1$, and so a contraction of a pair of vertices from L_i may create new red edges only to the remaining vertices of $L_{i-1} \cup L_i \cup L_{i+1}$.

Plane graphs; Left-aligned BFS trees. We will deal with *plane graphs*, which are planar graphs with a given (combinatorial) embedding in the plane, and one marked *outer face* (the remaining faces are then *bounded*). A plane graph is a *plane triangulation* if every face of its embedding is a triangle. Likewise, a plane graph is a *plane quadrangulation* if every face of its embedding is of length 4. It is easy to turn an embedding of any simple planar graph into a simple plane triangulation by adding vertices and incident edges into each non-triangular face. Furthermore, twin-width is non-increasing when taking induced subgraphs, and so it suffices to focus on plane triangulations in the proof of Theorem 1, and to similarly deal with plane quadrangulations in the proof of Theorem 2.

For algorithmic purposes, we represent a plane graph G in the standard combinatorial way – as a graph (the vertices and their adjacencies) with the counter-clockwise cyclic orders of the incident edges of each vertex, and we additionally mark the outer face of G.

In this planar setting, consider now a plane graph G, and a BFS tree T spanning G and rooted in a vertex r of the outer face of G, and picture (for clarity) the embedding G such that r is the vertex of G most at the top. For two adjacent vertices $u, v \in V(G)$, $\{u, v\} \in E(G)$, we say that u *is to the left of* v (wrt. T) if neither of u, v lies on the vertical path from r to the other, and the following holds; if r' is the least common ancestor of u and v in T and $P_{r',u}$ (resp., $P_{r',v}$) denote the vertical path from r' to u (resp., v), then the cycle $(P_{r',u} \cup P_{r',v}) + uv$ has the triple (r', u, v) in this counter-clockwise cyclic order.

A BFS tree T of G with the BFS layering $\mathcal{L} = (L_0, L_1, \ldots)$ is called *left-aligned* if there is no edge $f = uv$ of G such that, for some index i, $u \in L_{i-1}$ and $v \in L_i$, and u is to the left of v (an informal meaning is that one cannot choose another BFS tree of G which is "more to the left" of T in the geometric picture of G and T, such as by picking the edge uv instead of the parental edge of v in T).

▶ **Lemma 4.** *Given a simple plane graph G, and a vertex r on the outer face, there exists a left-aligned BFS tree of G and it can be found in linear time.*

Proof. For this proof, we have to extend the above relation of "being left of" to edges emanating from a common vertex of G. So, for an arbitrary BFS tree T of G and edges $f_1, f_2 \in E(G)$ incident to $v \in V(G)$, such that neither of f_1, f_2 is the parental edge of v in T, we write $f_1 \leq_l f_2$ if there exist adjacent vertices $u_1, u_2 \in V(G)$ such that u_1 is to the left of u_2, the least common ancestor of u_1 and u_2 in T is v and, for $i = 1, 2$, the edge f_i lies on the vertical path from u_i to v. Observe the following; if f_0 is the parental edge of v in T (or, in case of $v = r$, f_0 is a "dummy edge" pointing straight up from r), then $f_1 \leq_l f_2$ implies that the counter-clockwise cyclic order around v is (f_0, f_1, f_2). In particular, $\leq_l$ can be extended into a linear order on its domain.

We first run a basic linear-time BFS search from r to determine the BFS layering $\mathcal{L}$ of G. Then we start the construction of a left-aligned BFS tree $T \subseteq G$ from $T := \{r\}$, and we recursively (now in a "DFS manner") proceed as follows:

- Having reached a vertex $v \in V(T) \subseteq V(G)$ such that $v \in L_i$, and denoting by $X := (N_G(v) \cap L_{i+1}) \setminus V(T)$ all neighbours of v in L_{i+1} which are not in T yet, we add to T the nodes X and the edges from v to X.
- We order the vertices in X using the cyclic order of edges emanating from v to have it compatible with $\leq_l$ at v, and in this increasing order we recursively (depth-first, to be precise) call the procedure for them.

The result T is clearly a BFS tree of G. Assume, for a contradiction, that T is not left-aligned, and let $u_1 \in L_{i-1}$ and $u_2 \in L_i$ be a witness pair of it, where $\{u_1, u_2\} \in E(G)$ and u_1 is to the left of u_2. Let v be the least common ancestor of u_1 and u_2 in T, and let v_1 and v_2 be the children of v on the T-paths from v to u_1 and u_2, respectively. So, by the definition, $vv_1 \leq_l vv_2$ at v, and hence when v has been reached in the construction of T, its child v_1 has been taken for processing before the child v_2. Consequently, possibly deeper in the recursion, u_1 has been processed before the parent of u_2 and, in particular, the procedure has added the edge u_1u_2 into T, a contradiction to u_1 being to the left of u_2.

This recursive computation is finished in linear time, since every vertex of G is processed only in one branch of the recursion, and one recursive call takes time linear in the number of incident edges (to v). ◀

Notice that we have not assumed G to be a triangulation in the previous definition and in Lemma 4, which will be useful for the case of bipartite planar graphs.

Vertex levels in contraction sequences. We are going to work with contraction sequences which, preferably, preserve the BFS layers of $\mathcal{L}$ of connected G. However, we do *not always* preserve the layers, and so we need a notion which is related to the layers of $\mathcal{L}$, but it can differ from these layers when needed – informally, when this "causes no harm at all". For the graph G itself, we define $\lambda[G](v) = i$ if and only if $v \in L_i \in \mathcal{L}$. If G' is a trigraph along a contraction sequence of G, and a vertex $v' \in V(G')$ stems from a set $X \subseteq V(G)$ by (possible) contractions, then $\lambda[G'](v')$ equals the minimum i such that $L_i \cap X \neq \emptyset$. We say that $\lambda[G'](v')$ is the *level of* v' *in* G' along the considered contraction sequence of G, or simply the *level of* v' when the particular graph of a sequence is implicit. In other words, we can inductively say that if v'' of G'' results by the contraction of u' and v' of G', then $\lambda[G''](v'') := \min(\lambda[G'](u'), \lambda[G'](v'))$. If w' does not participate in a contraction along the subsequence from G' to G'', then $\lambda[G''](w') := \lambda[G'](w')$.

A *partial contraction sequence* of G is defined in the same way as a contraction sequence of G, except that it does not have to end with a single-vertex graph. A partial contraction sequence of G is *level-respecting* if every step contracts, in a trigraph G' along the sequence, only a pair $x, y \in V(G')$ such that the following inductively holds; the levels of x and y are the same, i.e. $\lambda[G'](y) = \lambda[G'](x)$, or all neighbours of y (red or black) in G' are on the same level as x is on, i.e. $\lambda[G'](z) = \lambda[G'](x)$ for all z such that $\{y, z\} \in E(G')$. (The conditions in the latter case, in particular, imply that $\lambda[G'](y) = \lambda[G'](x) + 1$; cf. Claim 5.)

Usefulness of level-respecting contraction sequences lies in the subsequent claim. Informally, we may say that our levels in G' behave analogously to the BFS layers of G; the levels form a layering (in the usual sense), albeit not ncessarily a BFS layering.

▷ Claim 5. Let a trigraph G' result from a level-respecting partial contraction sequence of a connected graph G. Then any vertex $z \in V(G')$ may have neighbours (red or black) only on the levels $\lambda[G](z) - 1$, $\lambda[G](z)$ and $\lambda[G](z) + 1$. Moreover, z must have some neighbour on the level $\lambda[G](z) - 1$.

Proof. We proceed easily by induction. The claim is trivial from the definition of a BFS layering when $G' = G$ and G is connected. Assume that $z \in V(G'')$ results from a contraction of a pair $x, y \in V(G')$, where $\lambda[G''](z) = \lambda[G'](x)$. Then $\lambda[G'](y) \in \{\lambda[G'](x), \lambda[G'](x) + 1\}$ by the definition of a level-respecting contraction and connectivity of G. So, there cannot be any neighbour of z on the levels lower than $\lambda[G''](z) - 1 = \lambda[G'](x) - 1$ from the induction. Regarding levels higher than $\lambda[G''](z) + 1$, they cannot host any neighbour of x in G' by the induction, and no neighbour of y as well by the definition of a level-respecting contraction (if $\lambda[G'](y) = \lambda[G'](x) + 1$). Lastly, since x has a neighbour on the level $\lambda[G''](z) - 1$ in G', so does z in G''. ◁

3 Proof of Theorem 1

3.1 Induction setup for a bounded region of the graph

Our main proof proceeds by induction on suitably defined subregions of the assumed plane triangulation G. In this subsection, we define the setup of this induction in Lemma 6, and show how it will imply the main result.

For a plane graph G and its cycle C, the *subgraph of G bounded by C*, denoted by G_C, is the subgraph of G formed by the vertices and edges of C and the vertices and edges of G drawn inside C – formally, in the region of the plane bounded by C and not containing the outer face. Let the vertices in the set $U := V(G_C) \setminus V(C)$ be called the *interior vertices* of C. We call a set $U_0 \subseteq U$ an *interior section of C* in G if all neighbours of vertices of U_0 belong to $U_0 \cup V(C)$ (in other words, U_0 is a collection of connected components of $G[U]$).

Consider a now fixed BFS tree T of G. Assume that a cycle C of G is formed as $C = (P_1 \cup P_2) + f$, where P_1 and P_2 are two T-vertical paths of length at least 1 with a common end $u \in V(P_1) \cap V(P_2)$ and $f \in E(G)$ is an edge joining the other ends v_1 of P_1 and v_2 of P_2. Observe that u is the (unique) vertex of G_C closest to the root r of T. Then we say that C is a *V-separator* in G with respect to implicit T ('V' as vertical), and we call u the *sink of C* and f the *lid of C*. If the vertices u, v_1, v_2 lie on C in this counter-clockwise order (equivalently, if v_1 is to the left of v_2 with respect to T), then we say that P_1 is the *left path* of C and P_2 is the *right path* of C (picture the sink at the top).

▶ **Lemma 6.** *Let G be a simple plane triangulation, and T be a left-aligned BFS tree of G rooted at a vertex $r \in V(G)$ of the outer triangular face and defining the initial levels $\lambda[G](\cdot)$. Assume that a cycle C of G is a V-separator of G, that G_C is the subgraph of G bounded by C, and u is the sink of C. Let the distance of u from the root r be ℓ, so $\lambda[G](u) = \ell$, and the maximum distance from a vertex of C to r be $m \geq \ell + 1$. Let $U \subseteq V(G_C)$ be an interior section of C in G, and denote by $W := V(G) \setminus (V(C) \cup U)$ the set of the "remaining" vertices.*

Then there exists a level-respecting partial contraction sequence of G which contracts only pairs of vertices that are in or stem from U, results in a trigraph G^*, and satisfies the following conditions for every trigraph G' along this sequence from G to G^*:

(I) *For $U' := V(G') \setminus (V(C) \cup W)$ (which are the vertices that are in or stem from U in G'), every vertex of U' in G' has red degree at most 8,*

(II) *every vertex of the left path of C has at most 5 red neighbours and every vertex of the right path of C has at most 3 red neighbours in U',*

(III) *the sink u of C has no red neighbour in U', and if the least level of a vertex of U in G is $k \geq \ell + 2$, then the vertices of C on levels up to $k - 2$ in G have no red neighbours in U' as well and each of the (two) vertices of C on the level $k - 1$ in G has at most 1 red neighbour in U', and*

(IV) *at the end of the partial contraction sequence, for the set $U^* := V(G^*) \setminus (V(C) \cup W)$ that stems from U in G^*, we have that if $z \in U^*$ is of level i, then $\ell < i \leq \max(m, \ell+2)$ and z is the only vertex in U^* of level i.*

Before proceeding further, we comment on two important things. First, we remark that, in Lemma 6, all vertices of U have the distance from r greater than ℓ, but on the other hand the distance from r to some vertices in U may be much larger than m (and our coming proof is aware of this possibility). Second, we observe that all vertices of U on level $\ell + 1$ must be adjacent to the sink u, since all other potential neighbours of them have the distance from r greater than ℓ. Consequently, contracting U on level $\ell + 1$ into one vertex within the claimed sequence indeed does not create a red edge to u, as long as we do not contract into it from higher levels (which we will explicitly avoid in the proof). We illustrate Lemma 6 in Figure 1.

We also observe that the assumptions and conditions of Lemma 6 directly imply some other properties useful for the coming proofs.

▷ Claim 7. *Respecting the notation and assumptions of Lemma 6, we also have that:*

(V) *Every red edge in G' has one end in U' and the other end in $U' \cup V(C)$,*

(VI) *if P_1 and P_2 are the left and right paths of C, respectively, and $v \in V(P_2)$ is of level j in G', then there is no edge of G' (red or black) from v to a vertex of $U' \cup (V(P_1) \setminus \{u\})$ of level $j - 1$ in G',*

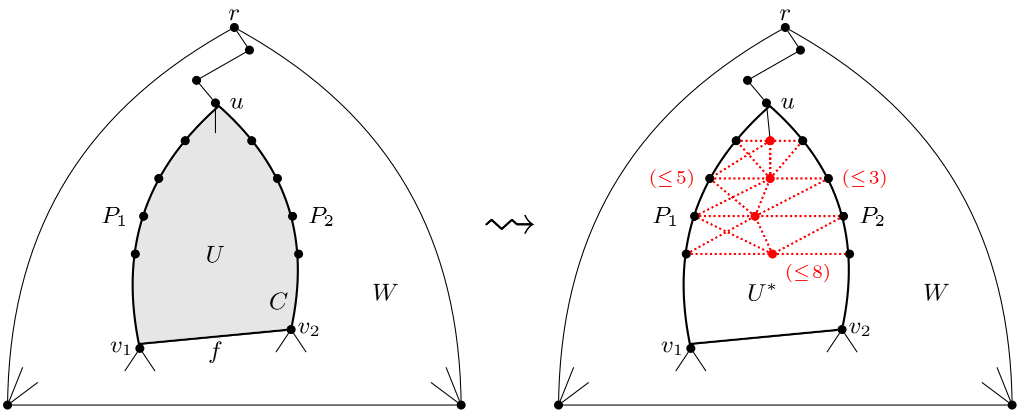

Figure 1 (left) The setup of Lemma 6, where P_1 and P_2 are the left and right paths of the chosen V-separator C. (right) The outcome of the claimed partial contraction sequence which contracts only vertices of U inside the shaded region from the left, and which maintains bounded red degrees in the region and on its boundary C. No other vertex than the sketched ones is affected by the contraction sequence. Not all depicted red edges do exist, and some of them may actually be black.

(VII) *at the end, that is, in G^*, every vertex of the left path of C has at most 3 red neighbours and every vertex of the right path of C has at most 2 red neighbours in U^*.*

Proof. Regarding (V), observe that since only vertices that stem from U participate in contractions, every red edge of G' must have an end in U'. Furthermore, since U is an interior section of C, no vertex of U is adjacent to a vertex of W in G, and hence no vertex of W is ever adjacent to a vertex being contracted in our sequence from G to G^*.

Concerning (VI), if v were adjacent to $x \in V(P_1)$ of level $j-1$, then this was already true in G; $\{x, v\} \in E(G)$. If v were adjacent to $x' \in U'$ of level $j-1$ in G', then, among the vertices of U contracted into x', there had to be $x \in U$ such that $\{x, v\} \in E(G)$. By the definition of a level-respecting sequence, possible vertices of level higher than $j-1$ contracted into x' cannot be adjacent to v of level j, and so $\lambda[G](x) = j-1$, too. Since, in both cases, such x would be to the left of v in G, this contradicts the assumption that T is left-aligned.

Finally, (VII) directly follows from Claim 5 and (IV) for the left path of C. For the right path we additionally apply (VI), which for $x \in V(P_2)$ of level j says that potential red neighbours of x are only on levels j and $j+1$. $\lhd$

We also show how Lemma 6 implies the first part of our main result:

Proof of Theorem 1 (the upper bound). We start with a given simple planar graph H, and extend any plane embedding of H into a simple plane triangulation G such that H is an induced subgraph of G. Then we choose a root r on the outer face of G and, for some left-aligned BFS tree of G rooted in r which exists by Lemma 4, the facial cycle C of the outer face incident to r, and $u = r$, we apply Lemma 6.

This way we get a partial contraction sequence from G to a trigraph G^* of maximum red degree 8 (along the sequence). Observe by (IV) that the set $U^* = V(G^*) \setminus V(C)$ contains only two vertices, on levels 1 and 2. In the final phase, we may hence pairwise contract the remaining vertices in an arbitrary order. The restriction of this whole contraction sequence of G to only $V(H)$ then certifies that the twin-width of H is at most 8. ◀

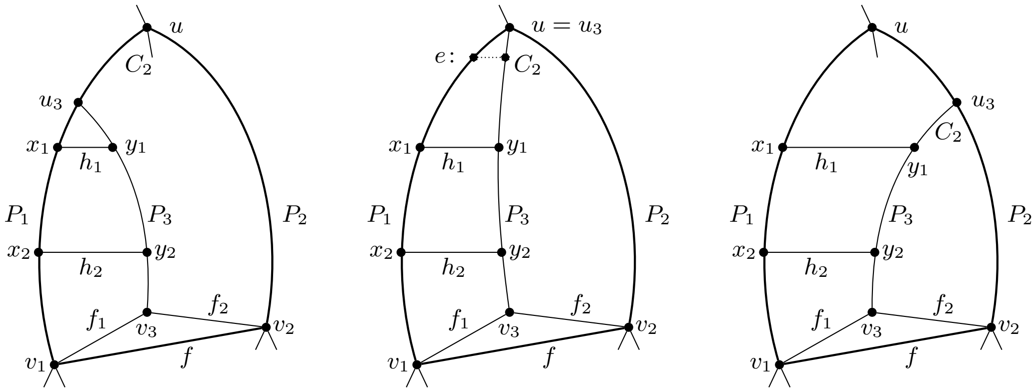

Figure 2 Three schematic cases of the vertical-horizontal division of G_C discussed in Section 3.2. The (possible) dotted horizontal edge e in the middle case is not counted as h_1 since no vertex of U is on a level above it in the picture.

3.2 Vertical-horizontal division into subregions

In order to apply induction in the proof of Lemma 6, we need to decompose the considered subgraph into suitable subregions, based on the plane drawing. Here we formulate the general decomposition step, while possible degenerate cases will be handled later in Section 3.3.

Let C be a V-separator in the plane triangulation G, formed by the left path P_1, the right path P_2 and the lid edge $f = \{v_1, v_2\}$ where v_i is an end of P_i. Let G_C be the subgraph bounded by C and $U \subseteq V(G_C)$ be an interior section of C in G. Moreover, assume that there exists a triangular face in G incident to f with the vertices v_1, v_2, v_3 where $v_3 \in U$, and that the T-vertical path from v_3 to the root r contains neither v_1 nor v_2. In particular, since T is left-aligned, we have that $\lambda[G](v_2) \leq \lambda[G](v_1)$ and $\lambda[G](v_3) \leq \lambda[G](v_1)$. Under these assumptions, we are going to define the *vertical-horizontal division* of G_C as follows.

Let $P \subseteq T$ be the vertical path connecting v_3 to the root r where, by $r \notin U$ and planarity of G, we have that P contains the sink u. Let $P_3 \subseteq P$ be the subpath of P from v_3 to the first vertex $u_3 \in V(P) \cap V(C)$ shared with the cycle C. We have $u_3 \neq v_3$. (It may be that $u_3 \in V(P_1)$ or $u_3 \in V(P_2)$ or even $u_3 = u$; see in Figure 2.) Let P_{31} denote the subpath of P from v_3 to the first intersection x with P_1 ($x \in \{u_3, u\}$), and P_{11} the subpath of P_1 from v_1 till x. Similarly, let P_{32} denote the subpath of P from v_3 to the first intersection y with P_2, and P_{22} be the subpath of P_2 from v_2 till y. Let $f_1 = \{v_1, v_3\}$ and $f_2 = \{v_2, v_3\}$. Observe that $P_{31}, P_{32} \subseteq G_C$, and that $C_1 := (P_{11} \cup P_{31}) + f_1$ and $C_2 := (P_{22} \cup P_{32}) + f_2$ are again V-separators in G, such that P_{31} is the right path of C_1 and P_{32} is the left path of C_2.

Furthermore, let $h_1, \ldots, h_a$, $a \geq 0$, be the collection of all horizontal edges of G_C such that, for $h_i = \{x_i, y_i\}$, we have $x_i \in V(P_1)$, $y_i \in V(P_3) \setminus V(C)$ and $\lambda[G](x_i) = \lambda[G](y_i)$, and that $\lambda[G](z) \leq \lambda[G](x_i) - 1$ holds for some $z \in U$ (the reason for this strange-looking restriction is in property (III) of Lemma 6). These edges $h_1, \ldots, h_a$ are ordered by their increasing level $\lambda[G](x_i)$. This is illustrated in Figure 2 (where the ordering of h_i's is top-down). For $i = 1, \ldots, a$, let $C_{1,i-1}$ denote the cycle passing through the sink of C_1 (which is u_3 or u) and formed by relevant subpaths of P_{11}, P_{31} and the edge h_i. Let $C_{1,a} = C_1$. Let $U_{1,0}$ denote the set of the interior vertices of $C_{1,0}$ in G, and for $i = 1, \ldots, a$, let $U_{1,i} := X \setminus U_{1,i-1}$ where X is the set of the interior vertices of $C_{1,i}$ in G. Let U_2 denote the set of the interior vertices of C_2 in G.

The system of the cycles $C_{1,0},\dots,C_{1,a},C_2$ and of the sets $U_{1,0},\dots,U_{1,a},U_2$ is called the *vertical-horizontal division* of G_C. The following is straightforward from the definition:

▷ Claim 8. For $i=0,1,\dots,a$, the cycle $C_{1,i}$ is a V-separator in G, and each vertex of $U_{1,i}$ has neighbours only in $U_{1,i}\cup V(C_{1,i})$. Hence, $U_{1,i}$ is an interior section of $C_{1,i}$. Consequently, for every $z\in U_{1,i}$ where $i\geq 1$, we have $\lambda[G](z)\geq\lambda[G](x_i)+1$ (where $\{x_i,y_i\}=h_i$ above).

The intended purpose of a vertical-horizontal division in the proof of Lemma 6 is to start the induction step, as precisely formulated in the next lemma with a straightforward proof:

▶ **Lemma 9.** *Assume the notation and assumptions of Lemma 6 for the graph G, cycle C and set U, and consider the vertical-horizontal division of the subgraph G_C as defined above; that is, the cycles $C_{1,0},\dots,C_{1,a},C_2$ and the sets $U_{1,0},\dots,U_{1,a},U_2$. Then the following hold:*
a) Each cycle $C^1\in\{C_{1,0},\dots,C_{1,a},C_2\}$ and the corresponding set $U^1\in\{U_{1,0},\ \dots,U_{1,a},U_2\}$ satisfy the assumptions of Lemma 6 (in the place of C and U).
b) Let $\tau_{1,i}$, $i=0,\dots,a$, denote the level-respecting partial contraction sequence of G claimed by Lemma 6 for the input as in (a) $C^1:=C_{1,i}$ and $U^1:=U_{1,i}$, and likewise, τ_2 be that for the input $C^1:=C_2$ and $U^1:=U_2$. Then the concatenated partial contraction sequence $\sigma_0:=\tau_2\cdot\tau_{1,0}\cdot\ldots\tau_{1,a}$, i.e., one starting with τ_2 and ending with $\tau_{1,a}$, again satisfies the properties (I), (II) *and* (III) *of Lemma 6.*

Proof. Part (a) immediately follows from the definition and Claim 8.

In part (b), we first argue that the concatenation σ_0 is well-founded; that contractions in one of the subsequences of σ_0 have no effect on vertices being contracted in another of the subsequences. This is since, by Claim 8, neighbours of contracted pairs of one subsequence are only in the interior section of the same subsequence or on the bounding cycles which together form $V(C)\cup V(P_3)$ (and the latter set is not participating in the contractions of σ_0).

Following on the previous argument, we have that the only vertices of G_C that may potentially receive red edges from more than one of the subseqeunces forming σ_0, are those of $V(C)\cup V(P_3)$. See Figure 2. For all other vertices of G_C, we have that (I) is true for them along whole σ_0 since it has been true for them along their subsequences of σ_0.

For a vertex $z\in V(P_3)\setminus V(C)$, we see that z belongs to C_2 and to each of $C_{1,i},\dots,C_{1,a}$ for some $i\in\{0,\dots,a\}$. However, even if $i\leq a-2$, vertices of $U_{1,i+2}$ cannot be neighbours of z in G_C due to a combination of Claim 5 and Claim 8. Therefore, z may have neighbours (and so can get red edges from by contractions) only in the interior sections U_2 and $U_{1,i}$, and possibly in $U_{1,i+1}$ if z is an end of the horizontal edge h_{i+1}. Recall also that z belongs to the right path of $C_{1,i}$. Along the sequence σ_0, but before $\tau_{1,i+1}$, the vertex z has red degree at most $5+3=8$ by (II) applied to U_2 and $U_{1,i}$. After $\tau_{1,i}$ is finished, z has red degree at most $3+2=5<8$ by (VII) of Claim 7. So, along the rest of σ_0, (I) stays true for z with red degree at most $5+1=6$ by (III) applied possibly to $U_{1,i+1}$.

For the vertex u itself, (III) is true automatically. For $z\in V(P_1)\setminus\{u\}$ and $i\in\{0,\dots,a\}$ being the least index such that $z\in V(C_{1,i})$, we get that the properties are true along $\tau_{1,i}$ and before by (II), and since $\tau_{1,i}$ ends, the vertex z has at most 3 red neighbours in $U_{1,i}$ by (VII). Additionally, z may get at most 1 red neighbour in $U_{1,i+1}$ (and none in $U_{1,i+2},\dots$) by (III), altogether at most 4, satisfying (II). In the special case of $z\in V(P_1)\setminus\{u\}$ covered by property (III) with respect to C, that is when all vertices belonging to the interior of C_1 are on levels higher than $\lambda[G](z)$, we get that this property is satisfied by (III) with respect to $C_{1,i}$, and there are no more red neighbours of z from elsewhere.

Finally, for $z\in V(P_2)\setminus\{u\}$, the conditions are simply true by (II) and possibly (III) for τ_2 and then along the whole sequence σ_0. ◀

3.3 Finishing the proof

Now we get to the core proof of Lemma 6 which will conclude our main result.

Proof of Lemma 6. We first resolve several special cases. If $U = \emptyset$, we are immediately done with the empty partial contraction sequence. So, assume $U \neq \emptyset$.

Recall the edge $f = \{v_1, v_2\} \in E(G)$ connecting the other ends of the left path P_1 and the right path P_2 of C. See again Figure 2. If v_1 has no neighbour in U, then $\{v_2, v_3\} \in E(G)$ where v_3 is the neighbour of v_1 on P_1. In such case, we simply apply Lemma 6 inductively to $P_1 - v_1$ and P_2, while the rest of the assumptions remain the same. The symmetric argument is applied when v_2 has no neighbour in U.

Otherwise, let $v_3 \in U$ be the vertex (unique in U) such that (v_1, v_2, v_3) bound a triangular face of G. Let $P \subseteq T$ be the vertical path connecting v_3 to the root r. If $v_1 \in V(P)$, then $P \supset P_1$ and we (similarly as above) apply Lemma 6 inductively to the V-separator $C^1 = P \cup P_2$ with the lid $\{v_2, v_3\}$, while the rest of the assumptions again remain the same. Note that in this case, $\lambda[G](v_1) = \lambda[G](v_3) - 1 = \lambda[G](v_2)$ since T is left-aligned. In the resulting trigraph G^1, we have the set $U^1 := V(G^1) \setminus (V(C^1) \cup W)$ that stems by contractions from the interior of C^1. There is no vertex in U^1 of level higher than $\lambda[G](v_3) \geq 2$ and at most one of level equal to $\lambda[G](v_3)$, by (IV) of Lemma 6. We contract the latter vertex with v_3, and then with the vertex of U^1 of the previous level $\lambda[G](v_1)$ unless v_1 is a neighbour of u (cf. the special case in (IV)). This clearly does not exceed red degree 8 there, and does not add new potential red neighbours to the vertices of C. Since (IV) is now satisfied, too, we are done. If $v_2 \in V(P)$, we solve the case similarly by induction applied to the V-separator $C^1 = P_1 \cup P$ with the lid $\{v_1, v_3\}$.

In all other cases, we have got a vertical-horizontal division of the subgraph G_C, with $P_3 \neq \{v_3, u\}$, with the horizontal edges $h_1, \ldots, h_a$, $h_i = \{x_i, y_i\}$, the cycles $C_{1,0}, \ldots, C_{1,a}, C_2$ and the interior sets $U_{1,0}, \ldots, U_{1,a}, U_2$, and we apply Lemma 9 to it. This way we get a level-respecting partial contraction sequence σ_0, which satisfies the properties (I), (II) and (III) of Lemma 6. Let G^0 denote the trigraph which results from G by σ_0, and let $U^0_{1,0}, \ldots, U^0_{1,a}, U^0_2$ denote the vertex sets of G^0 that stem from $U_{1,0}, \ldots, U_{1,a}, U_2$, respectively.

We first consider a subcase, that P_3 consists of a single edge $\{v_3, u_3\}$ and there is no vertex $z \in U$ in G such that $\lambda[G](z) \leq \lambda[G](u_3)$. This subcase has to be treated specially to fulfill (III) of Lemma 6. Then $a = 0$ in the vertical-horizontal division of G_C, and $\lambda[G](v_1) \leq \lambda[G](v_3) + 1 = \lambda[G](u_3) + 2$. Each of the sets $U^0_{1,0}$ and U^0_2 hence contains vertices at most on the levels $\lambda[G](v_3)$ and $\lambda[G](v_3) + 1$, by (IV) of Lemma 6.

We finish the desired partial contraction sequence from G^0 in this subcase by firstly contracting the two (if existing) vertices of $U^0_{1,0} \cup U^0_2$ on the level $\lambda[G](v_3) + 1$, and secondly by contracting each of the vertices of $U^0_{1,0} \cup U^0_2$ on the level $\lambda[G](v_3)$ with v_3. If $u_3 = u$ is the sink of C, then the only vertex on the level $\lambda[G](v_3) - 1 = \lambda[G](u)$ in G_C is u, and so every vertex of $U^0_{1,0} \cup U^0_2$ on the level $\lambda[G](v_3)$ must be adjacent to u (cf. Claim 5), and this is by a black edge due to an inductive invocation of (III). Therefore, the contractions into v_3 do not create a red edge to u. If $u_3 \neq u$, then let u'_3 denote the other vertex of $P_1 \cup P_2$ on the level $\lambda[G](u_3)$. Analogously to the previous case, one of the contractions into v_3 does not create a red edge to $\{u_3, u'_3\}$ and the other contraction can do so, but at most one red edge to each of u_3, u'_3. Therefore, (IV) is true here, and the remaining properties of Lemma 6 are fulfilled easily.

In the remaining cases, we possibly add the following bit in a sequence σ_1 after σ_0 (while this bit has not been possible in the special subcase above): If $u_3 \in V(P_1) \setminus \{u\}$ and u_3 is a neighbour of both x_1, y_1 (of h_1), then $U^0_{1,0}$ by (IV) consists of at most two vertices, which we

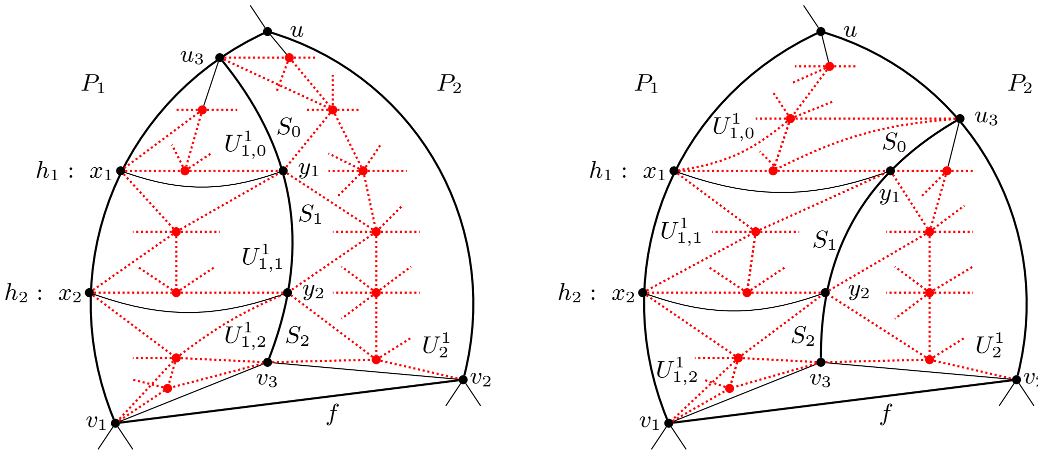

Figure 3 Proof of Lemma 6: a schematic picture of the situation after the parts of the depicted vertical-horizontal division of G_C have been recursively contracted (right before the σ_2-contractions start). The cases of u_3 on the left and right paths are not symmetric in general.

contract into one vertex in σ_1 – this move adds one red edge incident to u_3. Analogously, if $u_3 \in V(P_2) \setminus \{u\}$ and u_3 is a neighbour of both v_2, v_3 (of f_2), then we contract the at most two vertices of U_2^0 into one within σ_1. Although this contraction in σ_1 does not preserve levels, it is level-respecting by Claim 5 since $U_{1,0}$ and U_2 were interior sections of the triangles $C_{1,0}$ and C_2, respectively. In both cases, the added red edge incident to u_3 does not violate the properties of Lemma 6; this follows from the bounds in (VII) of Claim 7 which are by at least one lower than the bounds in (II) of Lemma 6, and property (III) is void for u_3 unless we have got the previous special subcase. Otherwise, we leave $\sigma_1 = \emptyset$.

After applying σ_1 to $U_{1,0}^0, \ldots, U_{1,a}^0, U_2^0$ in G^0, we get the trigraph G^1 and the sets $U_{1,0}^1, \ldots, U_{1,a}^1, U_2^1$ (which are identical to the former ones except possibly $U_{1,0}^0$ or U_2^0). See Figure 3.

In the next steps, we are going to define level-respecting partial contraction sequences $\sigma_{2,a}, \sigma_{2,a-1}, \ldots, \sigma_{2,0}$ which, when concatenated after $\sigma_0 \cdot \sigma_1$, give the desired outcome. If $a = 0$, the sequence $\sigma_{2,0}$ is going to contract the sets $U_{1,0}^1$ with $S_0 := V(P_3) \setminus \{u_3\}$ and $U_{2,0}^1 := U_2^1$. If $a > 0$, the sequence $\sigma_{2,a}$ is going to contract $U_{1,a}^1$ with the sets S_a and $U_{2,a}^1$, where $S_a \subseteq V(P_3) \setminus \{u_3\}$ and $U_{2,a}^1 \subseteq U_2$ are both the subsets of those vertices on levels greater than $\lambda[G](y_a)$. The sequence $\sigma_{2,i}$ for $0 \le i < a$ is going to contract $U_{1,i}^1$ with the sets S_i and $U_{2,i}^1$, where $S_i \subseteq V(P_3) \setminus \{u_3\}$ and $U_{2,i}^1 \subseteq U_2^1$ are the subsets of those vertices on levels greater (if $i \ge 1$) than $\lambda[G](y_i)$ and not greater than $\lambda[G](y_{i+1})$. Of course, some of these sets may be empty, and hence some contractions may not happen.

Specifically, for $i \in \{0, \ldots, a\}$ let $p = \max_{z \in C_{1,i}} \lambda[G](z)$ and $q = 1 + \min_{z \in C_{1,i}} \lambda[G](z)$. Observe that there is no vertex in $U_{1,i}^1 \cup U_{2,i}^1$ of level lower than q or greater than p. This follows from an inductive invocation of (IV) of Lemma 6, and from the sequence σ_1. So, the union $U_1^1 := U_{1,0}^1 \cup \ldots \cup U_{1,a}^1$ has at most one vertex on each level. Likewise, each of the sets $V(P_3)$ and U_2^1 has at most one vertex on each level. The sequence $\sigma_{2,i}$ first runs over $j = p, p-1, \ldots, q$ in this order, and contracts the pair of vertices of $S_i \cup U_{2,i}^1$ of the equal level j in G^1 (or nothing if there is at most one such vertex there). In its second round, $\sigma_{2,i}$ again runs over $j = p, p-1, \ldots, q$ in this order, and contracts the vertex of level j that stems from $S_i \cup U_{2,i}^1$ in the first round, with the vertex of $U_{1,i}^1$ of equal level j in G^1.

Let σ_2 be the concatenation of the described sequences, $\sigma_2 := \sigma_{2,a} \cdot \sigma_{2,a-1} \cdot \ldots \sigma_{2,0}$ in this order, and G^2 denote the trigraph which results from G^1 by applying σ_2. Let $U^2 := V(G^2) \setminus (V(C) \cup W)$ denote the contracted vertices in the interior of C in G^2. Then G^2 and U^2 satisfy property (IV) of Lemma 6 (in the place of G^* and U^*), which is immediate from the previous definition of σ_2. It thus remains to verify the properties (I), (II) and (III) of Lemma 6 along the sequence σ_2 from G^1 to G^2, that is, for every trigraph G' along σ_2.

Denote by $U' := V(G') \setminus (V(C) \cup W)$ all interior vertices of C in G', and by $U'' := U' \setminus V(G^1)$ the (new) interior vertices that stem by σ_2-contractions from G^1 to G', and recall (from Section 3.2) that $P_{31} \supseteq P_3$ is the right path of C_1 and $P_{32} \supseteq P_3$ is the left path of C_2 in G_C.

We start with verification of (III) which has already been in parts addressed above. Regarding the sink vertex u, it has got red edges neither from the sequence σ_0 by an inductive invocation of (III), nor from the sequence σ_1. The vertices of C on levels up to $k-2$ as in (III) do not have any neighbour in U' by Claim 5. Consider the vertices $z, z' \in V(C)$ on the level $k-1$ as in (III) (if $k \geq \ell+2$ there). If $\lambda[G](u_3) \geq k$, then no contraction on the level k happens within σ_2 (Figure 3), and so z and z' have each at most one red edge to U' by an inductive invocation of (III). Otherwise, up to symmetry, $z' = u_3$. Similarly as argued earlier in this proof, u_3 then has a black edge to $U^1_{1,0}$ (if $u_3 \in V(P_1)$) or to U^1_2 (if $u_3 \in V(P_2)$) in G^1, and so the contraction on the level k incident with this black edge does not create a new red edge to either of z, u_3. At the same time, each of z, u_3 has at most one red edge in G^1 by an inductive invocation of (III), and this stays true also (with the set U') during and after contractions on the level k within σ_2.

We move towards verification of (I). Let $z \in U'$ for the rest. If $z \in U^1_2$, then no σ_2-contraction has touched z so far. In this case z may have red edges to up to 3 vertices of $V(P_2) \cup U^1_2 \cup V(P_{32})$ of level $\lambda[G^1](z)-1$, to 2 vertices of $V(P_2) \cup V(P_{32})$ of level $\lambda[G^1](z)$, and to 2 vertices of $V(P_{32}) \cup U^1_2 \cup U''$ of level $\lambda[G^1](z)+1$, altogether at most 7. Note that there is no edge from z to the vertex of P_2 of level $\lambda[G^1](z)+1$ by (VI) of Claim 7. If $z \in V(P_3) \setminus \{u_3\}$, then similarly, z may have red edges to up to $2+2$ vertices of $U^1_1 \cup U^1_2$ on the levels $\lambda[G^1](z)-1$ and $\lambda[G^1](z)$, and to up to 2 vertices of $U^1_1 \cup U^1_2 \cup U''$ on the level $\lambda[G^1](z)+1$. The case of $z \in U^1_1$ (not-yet touched by a σ_2-contraction) is similarly easy.

Assume now that $z \in U''$ has been created in G' by a contraction of $z_2 \in U^1_2$ and $z_3 \in V(P_3) \setminus \{u_3\}$ (i.e., within the first round of some $\sigma_{2,i}$ above), but z is not contracted with a vertex of U^1_1 yet. Let $t \in V(P_1) \setminus V(P_3)$ denote the possible (unless equal to u_3) vertex of P_1 of level $\lambda[G^1](z_3)-1$. Then there is no edge in G^1 from t to z_2 by planarity, and no from t to z_3 by (VI) of Claim 7. The same applies to the possible vertex $t' \in U^1_1$ of level $\lambda[G^1](z_3)-1$. Consequently, z may have red edges to up to 3 vertices of $V(P_2) \cup U^1_2 \cup V(P_{32})$ of level $\lambda[G'](z)-1$, to 3 vertices of $V(P_2) \cup V(U^1_1) \cup V(P_1)$ of level $\lambda[G'](z)$, and to up to 3 vertices of $U'' \cup V(U^1_1) \cup V(P_1)$ of level $\lambda[G'](z)+1$, again using (VI). This sums to $3+3+3=9$, but we are going to show that this maximum of 9 cannot be achieved. Let $z_1 \in V(P_1)$ be such that $\lambda[G'](z_1) = \lambda[G'](z)$. If $\{z_1, z_3\} \notin E(G)$, then no red edge $\{z_1, z\}$ is created by the current contraction and the sum is at most 8, as needed. If $\{z_1, z_3\} = h_i \in E(G)$, then the sequence $\sigma_{2,i}$ has already contracted $S_i \cup U^1_{1,i}$ into U'', and so there are only 2 red neighbours of z in $U'' \cup V(P_1)$ on the level $\lambda[G'](z)+1$, again summing to at most 8. If $\{z_1, z_3\} \in E(G)$, but $\{z_1, z_3\}$ has not been chosen as any h_i in the vertical-horizontal division above, then there are no vertices in $U^1_2 \cup (V(P_3) \setminus \{u_3\})$ on the level $\lambda[G'](z)-1$, and so the sum is at most 8.

Assume that $z \in U''$ has already been created in G' by a contraction of all vertices in $U^1_2 \cup V(P_3) \cup U^1_1$ of the same level. Let this contraction be part of $\sigma_{2,i}$ for some $0 \leq i \leq a$. Then z may have red edges to up to 2 vertices of $V(P_2) \cup V(P_1)$ of level $\lambda[G'](z)$, to 2 vertices of $U'' \cup V(P_1)$ of level $\lambda[G'](z)+1$ (but not to $V(P_2)$ due to (VI)), and to vertices of $V(P_2) \cup U_2^1 \cup V(P_3) \cup V(U_1^1) \cup U'' \cup V(P_1) =: Y$ of level $\lambda[G'](z)-1$. However, at most 4 of the vertices of Y are potential red neighbours of z (so summing to at most 8), as we now show. If contractions on the level $\lambda[G'](z)-1$ are part of $\sigma_{2,i}$, too, then red neighbours of z in Y of level $\lambda[G'](z)-1$ actually belong to $V(P_2) \cup U'' \cup V(U_1^1) \cup V(P_1)$ with an upper bound of 4. Otherwise, if contractions on the level $\lambda[G'](z)-1$ are part of $\sigma_{2,i-1}$, then there is no edge from z to a vertex of U_{i-1}^1, or U_{i-1}^1 on the level $\lambda[G'](z)-1$ has already been contracted into U'', too. Then red neighbours of z in Y of level $\lambda[G'](z)-1$ belong to $V(P_2) \cup U_2^1 \cup V(P_3) \cup V(P_1)$ or to $V(P_2) \cup U'' \cup V(P_1)$, and we again get a bound of 4.

Finally, we want to verify (II) of Lemma 6. Consider $z \in V(P_2) \setminus \{u\}$. By the definition of σ_2, the vertex z may have at most one red neighbour of each level in U' (at any moment of σ_2). Then the bound of at most 3 red edges from z to U' follows immediately in view of Claim 5. Consider now $z \in V(P_1) \setminus \{u\}$, which is a bit more complicated case. On each of the levels $\lambda[G^1](z)-1$, $\lambda[G^1](z)$ and $\lambda[G^1](z)+1$ of U', there are clearly at most 2 red neighbours of z. Although, we now show that the maximum sum of 6 cannot be achieved. If $z = u_3$, then there is actually at most red neighbour of z on the level $\lambda[G^1](z)-1$. Otherwise, we denote the following vertices of G^1 of level $\lambda[G^1](z)+1$ by z_1, z_2, z_3 such that $z_1 \in U_1^1$, $z_2 \in U_2^1$, $z_3 \in V(P_3)$, and $\lambda[G^1](z_1) = \lambda[G^1](z_2) = \lambda[G^1](z_3) = \lambda[G^1](z)+1$. Then z_3 has no edge to $z \neq u_3$ by (VI) of Claim 7, and z_2 has no edge to z by planarity. If $z_3 \in U'$ (i.e., not contracted yet), then only z_1 may be a red neighbour of z. If z_2 and z_3 have already been contracted in G', but $z_1 \in U'$, then the new vertex again has no edge to z. Finally, if all of z_1, z_2, z_3 have been contracted in G', then U' has only (this) one vertex of level $\lambda[G^1](z)+1$. In any case, z has at most 5 red neighbours in U'.

We have verified all conditions of Lemma 6 for the partial contraction sequence $\sigma_0 \cdot \sigma_1 \cdot \sigma_2$, and so we can set $G^* := G^2$ and the proof is done. ◀

Proof of Theorem 1 (the algorithmic part). We can construct a simple plane triangulation $G \supseteq H$ in linear time using standard planarity algorithms, and then construct a left-aligned BFS tree $T \subseteq G$ again in linear time by Lemma 4. In the rest, we straightforwardly implement the recursive vertical-horizontal division of G as used in the proof of Lemma 6, and construct the contraction sequence of H on return from the recursive calls as defined in the proof. Note that we do not need at all to construct the intermediate trigraphs along the constructed contraction sequence, and so the construction of the sequence is very easy – each recursive call returns just a simple list of the vertices which stem from the recursive contractions, indexed by the levels. Then these (up to) two lists are easily in linear time "merged" together with the dividing path P_3, as specified by the proof of Lemma 6, into the resulting list of this call.

We may account total runtime in the "division part" of the algorithm to the edge(s) of v_3 into v_1 or v_2 and the edges of the path P_3 starting in v_3 in each call of the recursion, and these edges are not counted multiple times in different branches of the recursion. Likewise, runtime of the "merging" part of each recursive call can be counted to the individual steps of the resulting contraction sequence, which is of linear length. Hence, altogether, the algorithm runs in linear time. ◀

4 Proof of Theorem 2

On a high level, the proof will still proceed in the same way as in [16], and will prove the same bound. However, there are significant changes in the technical details, in which ideas from the previous section can save a lot of difficulties of the cumbersome proof from [16].

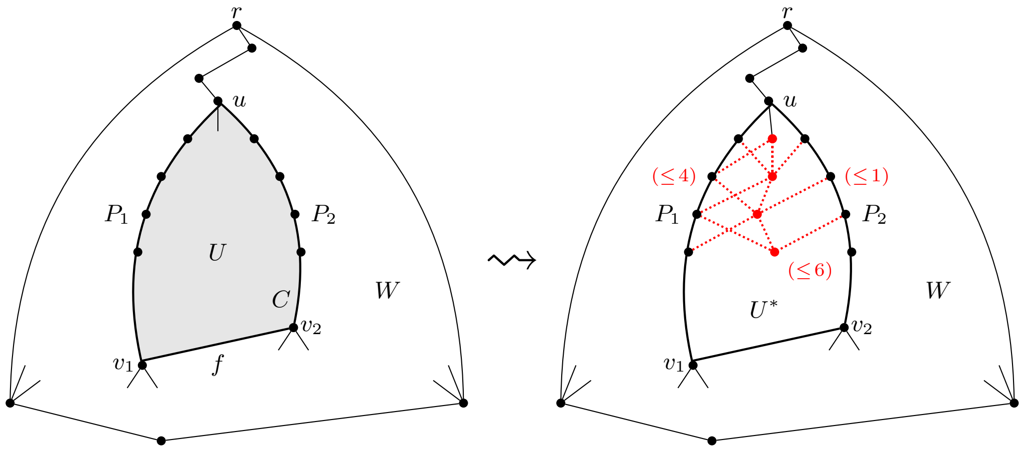

Figure 4 (left) The setup of Lemma 11, where P_1 and P_2 are the left and right paths of the chosen V-separator C. (right) The outcome of the claimed partial contraction sequence which contracts only vertices of U inside the shaded region from the left, and which maintains bounded red degrees in the region and on its boundary C.

In a nutshell, the bipartite case carries two major differences from the proof of the general planar case in Section 3:

- Since our graph is now bipartite, we will work with a plane quadrangulation (instead of a triangulation). However, with a suitable detailed analysis, it does not bring any significant new challenges to the proof.
- Since, again, our graph is bipartite, we immediately get that in any BFS layering, each layer is an independent set, and so we will never create a red edge inside the same layer. This is the crucial saving which allows us to derive a better upper bound on the red degree along the constructed sequence.

Before proceeding further, we need to adjust the concept of a level-respecting contraction sequence (beacuse of the fact that the definition from Section 2 possibly allowed to create new (red) edges inside the same level).

A partial contraction sequence of G is *bi-level-respecting* if every step contracts, in a trigraph G' along the sequence, only a pair $x, y \in V(G')$ such that the following inductively holds; the levels of x and y are the same, i.e. $\lambda[G'](y) = \lambda[G'](x)$, or all neighbours of y (red or black) in G' are on the level $\lambda[G'](x) + 1$, i.e. $\lambda[G'](z) = \lambda[G'](x) + 1$ is true for all z such that $\{y, z\} \in E(G')$. Again, we easily get by induction as in Claim 5:

▷ Claim 10. Let a trigraph G' result from a bi-level-respecting partial contraction sequence of a bipartite connected graph G. Then any vertex $z \in V(G')$ may have neighbours (red or black) only on the levels $\lambda[G](z) - 1$ and $\lambda[G](z) + 1$. In particular, the trigraph G' is again bipartite. Moreover, z must have some neighbour on the level $\lambda[G](z) - 1$.

Our proof is again by induction, precisely as set up in the following lemma. We illustrate this lemma in Figure 4. Before starting, note that a plane quadrangulation is always bipartite, and so it has all cycles (not only the faces) of length at least 4.

▶ **Lemma 11.*** *Let G be a simple plane quadrangulation, and T be a left-aligned BFS tree of G rooted at a vertex $r \in V(G)$ of the outer face and defining the initial levels $\lambda[G](\cdot)$. Assume that a cycle C of G is a V-separator of G, that G_C is the subgraph of G bounded by C, and u*

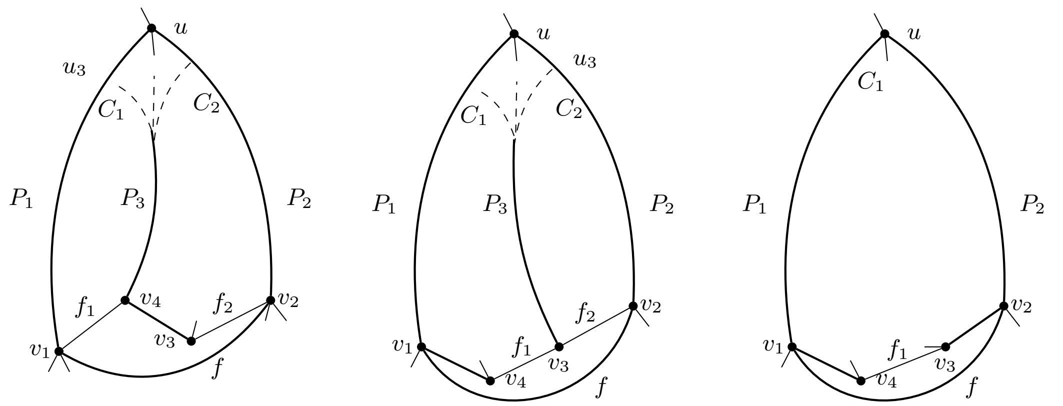

Figure 5 Three schematic cases of decomposing the drawing of G_C into subregions (bounded by C_1 and C_2) discussed in Lemma 13. They generally cover all possibilities in which $v_3, v_4 \notin V(P_1 \cup P_2)$. Vertical positions of the vertices of the 4-cycle $A = (v_1, v_2, v_3, v_4)$ outline their levels in G. Note that, in the right-most case, we have only one "nested" V-separator C_1, which is however not equal to $C = (P_1 \cup P_2) + f$, but instead C_1 passes through v_2, v_3, v_4, v_1 and has the lid edge $f_1 = \{v_3, v_4\}$.

is the sink of C. Let the distance of u from the root r be ℓ, so $\lambda[G](u) = \ell$, and the maximum distance from a vertex of C to r be $m \geq \ell + 2$. Let $U := V(G_C) \setminus V(C)$ be the interior vertices of C, and denote by $W := V(G) \setminus (V(C) \cup U)$ the set of the "remaining" vertices.

Then there exists a bi-level-respecting partial contraction sequence of G which contracts only pairs of vertices that are in or stem from U, results in a trigraph G^*, and satisfies the following conditions for every trigraph G' along this sequence from G to G^*:

- **(I)'** *For $U' := V(G') \setminus (V(C) \cup W)$ (which are the vertices that are in or stem from U in G'), every vertex of U' in G' has red degree at most 6,*
- **(II)'** *every vertex of the left path of C has at most 4 red neighbours and every vertex of the right path of C has at most 1 red neighbour in U',*
- **(III)'** *the sink u of C has no red neighbour in U',*
- **(IV)'** *if the next step of the sequence is going to contract a pair $x, y \in U'$ such that $\lambda[G'](y) > \lambda[G'](x)$, then y has no neighbour in the right path of C, and*
- **(V)'** *at the end of the partial contraction sequence, for the set $U^* := V(G^*) \setminus (V(C) \cup W)$ that stems from U in G^*, we have that if $z \in U^*$ is of level i, then $\ell < i \leq m + 1$ and z is the only vertex in U^* of level i.*

Lemma 11 already easily implies the first combinatorial part of Theorem 2, as we have seen with Theorem 1 in Section 3.1. Details of the algorithmic part again tightly follow the detailed proof steps which are present in the full paper.

In the proof of Lemma 11, we proceed analogously to Section 3. Namely, we start with a decomposition step analogous to Lemma 9 and illustrated in Figure 5. With a bit of technical work, we prove:

▶ **Lemma 13.*** *Assume the setting of Lemma 11, and with respect to it, let $U \neq \emptyset$ and $A \subseteq G_C$, $A = (v_1, v_2, v_3, v_4)$, denote the cycle bounding the 4-face incident to the lid edge f of the V-separator C and drawn in the closed disk of C.*

a) *There exists a vertical path $P_3 \subseteq G_C$ (internally disjoint from $C \cup A$ and possibly empty) such that the plane subgraph $C \cup A \cup P_3$ has two (if $P_3 = \emptyset$) or three distinct bounded faces, one of them being the face of A. The one or two bounded faces of $C \cup A \cup P_3$ other than that of A are bounded by cycles C_1 and C_2, where $v_1 \in V(C_1)$, and each of C_1 and C_2 (or just C_1 if $P_3 = \emptyset$) is again a V-separator whose lid edge is from $E(A) \setminus \{f\}$.*

b) *Assume now that $C \cup A \cup P_3$ has three bounded faces. Then $C_1 \cap C_2 = P_3$, and the sinks of C_1 and C_2 are in some order the sink u of C and the end u_3 of P_3 not in A (which may be the same vertex). If τ_i, $i = 1, 2$, is the bi-level-respecting partial contraction sequence of G obtained by inductively applying Lemma 11 to the cycle C_i, then the concatenation $\sigma_0 := \tau_2 \cdot \tau_1$ of these two sequences is a bi-level-respecting partial contraction sequence of G which satisfies the properties* (I)*' to* (IV)*' of Lemma 11.*

Then, we follow on the partial contraction sequence of Lemma 13 analogously to the proof in Section 3.3, albeit with slightly simpler arguments thanks to a simpler decomposition step with only at most two subregions. In this way we finish both Lemma 11 and Theorem 2, and the details are now left for the full paper [17].

5 Concluding Remarks

We have further improved by one the previous best upper bound [16] on the twin-width of planar graphs. This seemingly small improvement has required a careful reconsideration of the previous method and several new ideas, and although our new approach has simplified some cumbersome technical details in [16], new technical difficulties emerged which makes some parts of the proof again quite technical. This is probably to be expected since we are now very close to the currently best lower bound of 7 on the twin-width of planar graphs [19].

To recapitulate the fine improvements leading to the upper bound of 8 on the twin-width compared to previous larger bounds in [12, 18] and [3]; we communicate that the biggest (numerical) jump comes from the use of a specially tailored BFS-based decomposition formulated in Section 3.2, but we regard as the most important contribution in the quest the use of a left-aligned BFS tree (Section 2), which essentially “slashes down” additional up to three possible red neighbours from the analysis in the proof of Lemma 6. While both previous improvements have been introduced already in [16], the use of the “horizontal items” in the vertical-horizontal division of Section 3.2 then gives a final touch improving the bound to 8 (while 9 seemed to be unbeatable without this final trick).

Related to the twin-width is the notion of reduced bandwidth [12] which, informally stating, requires the subgraph induced by the red edges (along the sequence) to not only have bounded degrees, but also bounded bandwidth. Strictly speaking, as our construction of the contraction sequence creates arbitrarily large “red grids” in some cases, it does not directly imply any constant upper bound on the reduced bandwidth of planar graphs. However, a simple modification of the construction (informally, delaying contractions that would create red edges to the vertices of P_2 as in Figure 4) can easily bring a reasonable two-digit upper bound on the reduced bandwidth of planar graphs, which can possibly be further tightened with a specialized refined argument.

Besides the core question of the maximum twin-width of planar graphs, one may also reconsider the fact that our proof method is (distantly) based on the proof of the product structure of planar graphs [14] and ask whether we could possibly improve the product structure over the currently best variant in [21]. Unfortunately, our recursive decomposition of planar graphs is very tailored to the purpose of proving a good upper bound on the twin-width and it currently does not seem to yield an improvement in the planar product struture, or in the maximum queue number of planar graphs. This direction, however, is the subject of our ongoing research.

In the end we would like to dwell on the very idea of left-aligned BFS trees from Section 2. This seems like a quite general idea about planar graphs, related to other specialized BFS- and DFS-search routines in the algorithmic world, but we have not found this exact idea anywhere in the published literature (the existing related concepts we are aware of do not feature the BFS property). We believe that this new idea could possibly find its use in other problems regarding planar graphs and drawings.

To finally conclude, the problem to determine the exact maximum value of the twin-width over all planar graphs is still open, but our continuing research suggests that the value of 7 is much more likely (than 8) to be the right answer. Likewise, the problem to determine the exact maximum value of the twin-width over bipartite planar graphs is open, and we cannot now decide whether the value of 6 is the right maximum value over bipartite planar graphs, or whether the upper bound may possibly be 5 (while an upper bound lower than 5 is not likely since a bipartite construction analogous to [19] seems to exclude it, but we are not aware of this claim being written up as a formal statement).

References

1 Jungho Ahn, Kevin Hendrey, Donggyu Kim, and Sang-il Oum. Bounds for the twin-width of graphs. *CoRR*, abs/2110.03957, 2021. `arXiv:2110.03957`.

2 Jakub Balabán and Petr Hliněný. Twin-width is linear in the poset width. In *IPEC*, volume 214 of *LIPIcs*, pages 6:1–6:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

3 Michael A. Bekos, Giordano Da Lozzo, Petr Hliněný, and Michael Kaufmann. Graph product structure for h-framed graphs. *CoRR*, abs/2204.11495v1, 2022. `arXiv:2204.11495v1`.

4 Pierre Bergé, Édouard Bonnet, and Hugues Déprés. Deciding twin-width at most 4 is NP-complete. In *ICALP*, volume 229 of *LIPIcs*, pages 18:1–18:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

5 Édouard Bonnet, Dibyayan Chakraborty, Eun Jung Kim, Noleen Köhler, Raul Lopes, and Stéphan Thomassé. Twin-width VIII: delineation and win-wins. In *IPEC*, volume 249 of *LIPIcs*, pages 9:1–9:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

6 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. In *SODA*, pages 1977–1996. SIAM, 2021.

7 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: max independent set, min dominating set, and coloring. In *ICALP*, volume 198 of *LIPIcs*, pages 35:1–35:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

8 Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Torunczyk. Twin-width IV: ordered graphs and matrices. In *STOC*, pages 924–937. ACM, 2022.

9 Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, and Stéphan Thomassé. Twin-width VI: the lens of contraction sequences. In *SODA*, pages 1036–1056. SIAM, 2022.

10 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. In *FOCS*, pages 601–612. IEEE, 2020.

11 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. *J. ACM*, 69(1):3:1–3:46, 2022.

12 Édouard Bonnet, O-joung Kwon, and David R. Wood. Reduced bandwidth: a qualitative strengthening of twin-width in minor-closed classes (and beyond). *CoRR*, abs/2202.11858, 2022. `arXiv:2202.11858`.

13 Édouard Bonnet, Jaroslav Nesetril, Patrice Ossona de Mendez, Sebastian Siebertz, and Stéphan Thomassé. Twin-width and permutations. *CoRR*, abs/2102.06880, 2021. `arXiv:2102.06880`.

14 Vida Dujmovic, Gwenaël Joret, Piotr Micek, Pat Morin, Torsten Ueckerdt, and David R. Wood. Planar graphs have bounded queue-number. *J. ACM*, 67(4):22:1–22:38, 2020. `doi:10.1145/3385731`.

15 Jakub Gajarský, Michal Pilipczuk, Wojciech Przybyszewski, and Szymon Torunczyk. Twin-width and types. In *ICALP*, volume 229 of *LIPIcs*, pages 123:1–123:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

16 Petr Hliněný. Twin-width of planar graphs is at most 9, and at most 6 when bipartite planar. *CoRR*, abs/2205.05378, 2022. `arXiv:2205.05378`.

17 Petr Hliněný and Jan Jedelský. Twin-width of planar graphs is at most 8, and at most 6 when bipartite planar. *CoRR*, abs/2210.08620, 2022. `arXiv:2210.08620`.

18 Hugo Jacob and Marcin Pilipczuk. Bounding twin-width for bounded-treewidth graphs, planar graphs, and bipartite graphs. In *WG*, volume 13453 of *Lecture Notes in Computer Science*, pages 287–299. Springer, 2022.

19 Daniel Král and Ander Lamaison. Planar graph with twin-width seven. *CoRR*, abs/2209.11537, 2022. `arXiv:2209.11537`.

20 Michal Pilipczuk, Marek Sokolowski, and Anna Zych-Pawlewicz. Compact representation for matrices of bounded twin-width. In *STACS*, volume 219 of *LIPIcs*, pages 52:1–52:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

21 Torsten Ueckerdt, David R. Wood, and Wendy Yi. An improved planar graph product structure theorem. *CoRR*, abs/2108.00198, 2021. `arXiv:2108.00198`.