ON 2-COMPLEXES EMBEDDABLE IN 4-SPACE, AND THE EXCLUDED MINORS OF THEIR UNDERLYING GRAPHS

A. GEORGAKOPOULOS AND M. WINTER

Abstract. In 2006 van der Holst introduced a new 4D analogue of graph planarity: a graph is *4-flat* if every 2-dimensional CW complex on the graph embeds in $\mathbb{R}^4$. He showed that this class is minor-closed and stated a number of conjectures about the excluded minors, operations that preserve 4-flatness, and characterizations using the Colin de Verdiére graph invariant.

In this article we first study operations on 2-dimensional CW complexes that preserve embeddability in $\mathbb{R}^4$, such as joining and cloning 2-cells, as well as ΔY-transformations. We also construct a CW complex for which YΔ-transformations do not preserve embeddability. We then conclude a number of results on 4-flat graphs and answer several of van der Holst's conjectures: we show that 4-flat graphs are closed under cloning of edge and ΔY-transformations, we establish several non-trivially equivalent definitions of 4-flatness, and we give streamlined proofs for several properties of 4-flat graphs. Eventually we verify that all graphs of the Heawood family are indeed excluded minors for the class of 4-flat graphs.

1. Introduction

The study of planar graphs marks the starting point of both topological graph theory and graph minor theory. A variety of concepts have since been devised with the goal of capturing a higher-dimension analogue of planarity. Linkless graphs (or flat graphs), for example, are defined by the existence of particular embeddings into 3-space, and can likewise be characterized by a short list of excluded minors [13]. In 2006 van der Holst introduced a natural 4D analogue [17]: a graph G is *4-flat* if every 2-dimensional regular CW complex on G can be (piecewise linearly) embedded into $\mathbb{R}^4$. Van der Holst's work contains a number of results and plausible conjectures that paint the picture of 4-flat graphs as the "correct" continuation from planar and linkless to 4-space. In particular, all planar and linkless graphs are 4-flat.

At the same time the study of 4-flat graphs comes with a number of intricacies typical for questions in low-dimensional topology. The embedding problem $2 \to 4$ (*i.e.,* deciding embeddability of 2-dimensional CW complexes in 4-space), which seems required for deciding 4-flatness, is among the least understood. A fundamental homological obstruction to embeddability $n \to 2n$, the van Kampen obstruction, fails to be sufficient precisely for $n = 2$ [6], and it is an open question whether the decision problem $2 \to 4$ is even decidable [11]. As a consequence, many natural questions about 4-flat graphs are widely open.

In this article we study operations on 2-dimensional CW complexes that preserve embeddability in 4-space and make a number of conclusions about 4-flat graphs. We answer several conjectures of van der Holst regarding their excluded minors and on operations that preserve 4-flatness. Our tools allow us to give streamlined proofs for several known results.

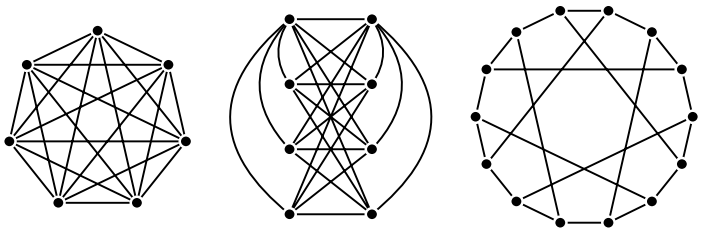

Figure 1. Three graphs from the Heawood family: K_7, $K_{3,3,1,1}$ and *the* Heawood graph. All three are not 4-flat.

1.1. **Background on 4-flat graphs.** Most results mentioned below were obtained by van der Holst in [17]. Some later work in [19] and [9] added to this.

The 4-flat graphs form a minor-closed graph family and are therefore characterized by a finite list of excluded minors (by the Robertson-Seymour theorem). The complete list of excluded minors is unknown, though plausible conjectures have been proposed: the *Heawood family* consists of the 78 graphs built from K_7 and $K_{3,3,1,1}$ by repeated application of ΔY- and YΔ-transformations. The members of the Heawood family are collectively known as "Heawood graphs", though not to be confused with *the* Heawood graph (Figure 1; right), which is however a member of this family. Van der Holst showed that the Heawood graphs are not 4-flat and suggested:

Conjecture 1.1 ([17])**.** *The excluded minors for the class of 4-flat graphs are exactly the 78 graphs of the Heawood family.*

Note the structural similarity of the Heawood family to the Kuratowski graphs (K_5 and $K_{3,3}$; the excluded minors for planarity) and the Petersen family (generated from K_6 and $K_{3,3,1}$; the excluded minors for linkless graphs).

Van der Holst proved that both K_7 and $K_{3,3,1,1}$ are excluded minors, though left this open for the other graphs of the Heawood family. We shall fill this gap in Section 6. Van der Holst proposed another concise argument for proving this, though based on two further conjectures:

Conjecture 1.2 ([17, Conjecture 1])**.** *Cloning an edge of a graph (i.e., replacing it by two parallel edge) preserves 4-flatness.*

Conjecture 1.3 ([17, Conjecture 2])**.** ΔY- *and* YΔ*-transformations of graphs preserve 4-flatness (cf. Figure 3).*

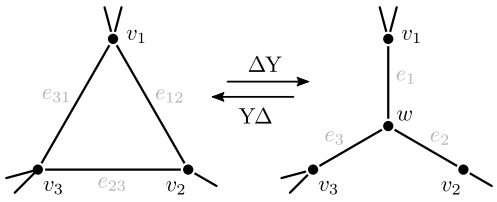

Figure 2. Visualization of ΔY- and YΔ-transformations.

We later confirm Conjecture 1.2 as well as the ΔY-part of Conjecture 1.3. We also provide constructions for complexes and graphs that highlight the difficulties in proving the YΔ-part of the conjecture.

Another striking analogy to planar and linkless graphs might exist in connection to the Colin de Verdiére graph invariant $\mu(G)$ [18]: recall that a graph G is planar if and only $\mu(G) \leq 3$, and is linkless if and only if $\mu(G) \leq 4$. Van der Holst conjectured:

Conjecture 1.4 ([17, Conjecture 3])**.** *A graph G is 4-flat if and only if it has Colin de Verdiére invariant $\mu(G) \leq 5$.*

We list some evidence in favor of this conjecture: all Heawood graphs have $\mu = 6$. It moreover holds that G is linkless if and only if the suspension $G * K_1$ (adding a new dominating vertex) is a 4-flat graph (one direction was proven in [17, Theorem 2], we prove the other direction in Theorem 5.9). It follows that $G * K_1$ is 4-flat if and only if $\mu(G * K_1) \leq 5$.

An alternative definition of 4-flat graphs based on "almost embeddings" as well as generalizations to even higher dimensions are introduced in [19]. Based on this the authors define a graph invariant $\sigma(G)$ and show that it agrees with $\mu(G)$ on linkless graphs, though diverges at sufficiently large values. It is unknown whether $\sigma(G) \leq 5$ is equivalent to being 4-flat. It holds $\sigma(G) \leq \mu(G)$ [9].

1.2. **Overview and results.** The results of this paper can be divided into results on general 2-dimensional CW complexes and result on 4-flat graphs.

Section 3 explores operations on general 2-dimensional CW complexes, and especially on when those preserve embeddability into $\mathbb{R}^4$.

Theorem. *Embeddable complexes are closed under any of the following operations:*

(*i*) *joining 2-cells at vertices and edges (Lemma 3.1).*
(*ii*) *cloning 2-cells (Lemma 3.2).*
(*iii*) *collapsing 2-cells (Lemma 3.9).*
(*iv*) *cloning edges (Corollary 3.11).*

We also obtain further results on rerouting 2-cells (Lemma 3.7), stellifying cycles (Lemma 3.15) (including the important special case of ΔY-transformations; Corollary 3.16), and merging parallel edges (Corollary 3.13). These operations preserve embeddability only under additional assumptions and so the precise statements requires more care.

As an application of these operations we proof the following:

Theorem 3.5. *In a 2-dimensional CW complex X, if there are two vertices $v, w \in X$ so that each 2-cell $c \subseteq X$ is incident to at least one of them, then X can be embedded in $\mathbb{R}^4$.*

To appreciate the strength of Theorem 3.5, we remark that it is non-trivial to prove that X is 4-embeddable even if its 1-skeleton has only a single vertex; see [4], where this statement is deduced from a difficult theorem of Stallings. Moreover, Theorem 3.5 becomes false if we replace v, w by a triple of vertices; a counterexample is the triple cone over any non-planar graph, as was proven by Grünbaum [8].

At this point we have the tools to verify van der Holst's Conjecture 1.2 (Corollary 3.12) and the ΔY-part of Conjecture 1.3 (Corollary 3.18) on 4-flat graphs.

Section 4 explores the intricacies of the inverse operations – reverse stellification, and in particular, YΔ-transformations. We show that the YΔ-part of Conjecture 1.3 is right at the boundary of what can be true: stellification at cycles of length $\ell \geq 4$

does not preserve 4-flatness (Example 4.2); and we construct an embeddable complex (based on the Freedman-Krushkal-Teichner complex) whose YΔ-transformation is not embeddable (Section 4.1).

From Section 5 on we explore the implications of our results to the theory of 4-flat graphs and their full complexes. Together with earlier results from Section 3, the following operations are shown to preserve 4-flatness:

Theorem. *4-flat graphs are closed under any of the following operations:*

(*i*) *cloning edges (Corollary 3.12).*
(*ii*) *ΔY-transformations(Corollary 3.18).*
(*iii*) *3-clique sums (Lemma 5.14).*

Other results obtained in Section 5 include:

- we consider alternative notions of 4-flatness: instead of *every regular* complex on G being embeddable, we require that *every* (not necessarily regular) complex, or *every induced* complex (*i.e.,* 2-cells only along induced cycles) is embeddable. We show that, surprisingly, both modifications give rise to the same notion of 4-flatness, despite being seemingly stronger and weaker respectively (Theorem 5.1).
- we show that 4-flat graphs are "locally linkless" (Lemma 5.7).
- we show that a suspension $G * K_1$ (adding to G a new dominating vertex) yields a 4-flat graph if and only if G is linkless (Theorem 5.9).
- we show that there is no simple inclusion relation between 4-flat and knotless graph: neither every 4-flat graph is knotless, nor the other way around (Section 5.5).

Eventually, in Section 6, we verify one direction of Conjecture 1.1:

Theorem 6.1. *All graphs of the Heawood family are excluded minors for the class of 4-flat graphs.*

2. Basic terminology and examples

2.1. **Graphs, complexes and embeddings.** Throughout the article G denotes a finite (multi-)graph, potentially with loops and multi-edges. We consider graphs as 1-dimensional CW complexes, *i.e.,* as topological spaces with 0-cells (the vertices) and 1-cells (the edges).

We write X to denote a *2-dimensional CW complex* (or just *complex* for short). A complex over G is obtained by attaching disks $\mathbb{D}^2$ (or *2-cells*) along closed walks in G. Each 2-cell $c \subset X$ is defined via its *attachment map* $\partial c\colon \partial\mathbb{D}^2 \to G$. If all 2-cells are attached along cycles (*i.e.,* closed walks without self-intersections) then X is said to be *regular*. Given a complex X, the underlying graph (or *1-skeleton*) is denoted G_X. For a cell $c \subseteq X$, we write $X \setminus c$ for the set-theoretic difference, and $X - c$ for the subcomplex with this cell removed (but its boundary ∂c left intact).

All subsets of $\mathbb{R}^d$, all embeddings into $\mathbb{R}^d$ and homeomorphisms between such sets are assumed to be piecewise linear (or PL). We follow piecewise linear topology as developed in [14].

All embeddings $\phi\colon X \to \mathbb{R}^d$ in this article are (if not stated otherwise) into $\mathbb{R}^4$. We write X^ϕ instead of $\phi(X)$ to denote the image of an embedding. A complex is *embeddable* (or *4-embeddable*) if there exists an embedding $\phi\colon X \to \mathbb{R}^4$. Note that all 2-complexes embed in $\mathbb{R}^5$, though not necessarily in $\mathbb{R}^4$. The classical example for

a non-embeddable complex is the triple cone over K_5 [8]. Let $\mathcal{K}_n$ be the complex obtained from the complete graph K_n by attaching a 2-cell along each triangle. Then $\mathcal{K}_7$ provides another example of a non-embeddable complex ([10, Theorem 5.1.1]). Removing a single 2-cell from $\mathcal{K}_7$ yields a complex $\mathcal{K}_7 - \Delta$ that is actually embeddable, and will serve a running (counter)example:

Example 2.1 ($\mathcal{K}_7 - \Delta$)**.** We have $\mathcal{K}_4 \simeq \mathbb{S}^2$ and $K_3 \simeq \mathbb{S}^1$. The join of two spheres is a sphere: $K_3 \star \mathcal{K}_4 \simeq \mathbb{S}^1 \star \mathbb{S}^2 \simeq \mathbb{S}^4$. Since $\mathcal{K}_7 - \Delta \subset K_3 \star \mathcal{K}_4$ (the missing 2-cell Δ corresponds precisely to the missing "filling" of K_3), this provides an embedding of $\mathcal{K}_7 - \Delta$ into $\mathbb{S}^4$.

2.2. **Full complexes and 4-flat graphs.** Throughout the text we use the following definition for 4-flat graphs:

Definition 2.2 (full complex, 4-flat graph)**.**

(*i*) The *full complex* $X(G)$ of G is the CW complex obtained by attaching a 2-cell along each cycle of G.

(*ii*) A graph is *4-flat* if its full complex is embeddable.

The full complex is clearly a regular complex. Since every regular complex is a subcomplex of $X(G)$, this is equivalent to the definition given in the introduction.

Some relevant examples of graphs that are not 4-flat follow from previous discussions: since $\mathcal{K}_7$ is not embeddable, K_7 is not 4-flat. Since the triple cone over K_5 is not embeddable, $K_5 * \overline{K}_3$ is not 4-flat. Since K_7 is known to be an excluded minor, $K_7 - e$ (K_7 with an edge removed) is 4-flat.

2.3. **Contraction.** Give an embeddable complex X and a collapsible subset $c \subseteq X$ (for example, a cell), the quotient complex X/c is embeddable as well. This fact is often used without much elaboration and can be justified roughly as follows: choose a small neighborhood B of c^ϕ and replace its interior with a cone over $\partial B \cap X^\phi$. We briefly remind of the technicalities involved in this argument:

Lemma 2.3. *If X is embeddable and $c \subseteq X$ is collapsible (e.g. because c is a cell of X) then X/c is embeddable as well.*

The proof requires the theory of regular neighborhoods (*e.g.* [14, Chapter 3]).

Proof sketch. Fix a triangulation of $\mathbb{R}^d$ with simplicial sub-complexes that triangulate X^ϕ and c^ϕ respectively (which exists by [14, Addendum 2.12]). Let $B \subset \mathbb{R}^4$ be a regular neighborhood of c^ϕ *w.r.t.* this triangulation (*e.g.* as defined in [14, Chapter 3]). Since c^ϕ is collapsible, B is a ball by [14, Corollary 3.27]. Then there exists a homeomorphism $\psi\colon B \xrightarrow{\sim} K$ with a convex set $K \subset \mathbb{R}^4$. Let C be the cone over $\psi(\partial B \cap X^\phi) \subset \partial K$ with apex at some interior point of K. Remove $X^\phi \cap B$ and replace it with $\psi^{-1}(C)$. This yields an embedding of X/c. □

3. Operations that preserve embeddability

In this section we study operations on general (2-dimensional) CW complexes and circumstances under which they preserve embeddability. The eventual goal is to develop a set of flexible tools to build and modify full complexes of 4-flat graphs. As an aside we obtain interesting results and conjectures for general complexes.

In order to be concise, we will often conflate an attachment map $\partial c\colon \partial\mathbb{D}^2 \to G_X$ with its image, considering ∂c as both a parametrized curve and a subset of X. This

picture is not accurate if ∂c is not injective, though we trust that the simplification will be understood as such. In case of doubt we ask the reader to first "regularize" the complex by subdividing 2-cells (ideally without introducing new vertices) until the complex is regular and all attachment maps are injective.

As the running (counter)example of this section we use the complex $\mathcal{K}_7$ that is obtained from K_7 by attaching a 2-cell along each triangle. Its vertices we denote by $x_1, ..., x_7$. The complex obtained by removing the triangle along $x_1x_2x_3$ we denote by $\mathcal{K}_7 - \Delta$. Recall that $\mathcal{K}_7$ does not embed, whereas $\mathcal{K}_7 - \Delta$ does (*cf.* Example 2.1).

3.1. **Joining 2-cells.** Let $c_1, c_2 \subseteq X$ be two distinct 2-cells and $\gamma \subset \partial c_1 \cap \partial c_2$ a path (potentially of length zero) in their shared boundary. *Joining* c_1 and c_2 at γ means to replace them by a new 2-cell c whose attachment map is a concatenation of $\partial c_1 - \gamma$ and $\partial c_2 - \gamma$.

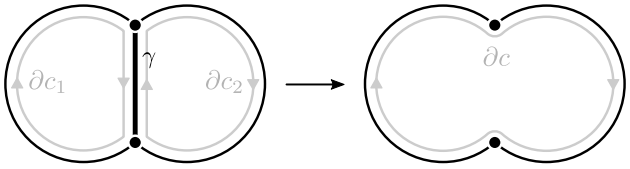

Figure 3. Joining 2-cells across a path γ.

Lemma 3.1. *Joining 2-cells at a path γ of lenght $\ell \leq 1$ (i.e., at a vertex or edge) preserves embeddability.*

Proof. Fix an embedding $\phi\colon X \to \mathbb{R}^4$. Choose a point $x \in \gamma^\phi$ and a link neighborhood $B \subset \mathbb{R}^4$ of x. The idea is to delete $B \cap c_1^\phi$ and $B \cap c_2^\phi$ and to fill in the resulting holes with suitable new patches that join the 2-cells into a single one. To define the patches, let $\rho_i := \partial B \cap c_i^\phi$ be the paths in the link that correspond to the original 2-cells. We proceed depending on the length of γ.

Case $\ell = 0$: There exists a square $[0,1]^2 \simeq Q \subseteq \partial B$ with two opposite edges attached along ρ_1 and ρ_2 respectively, disjoint from the rest of the link, and potentially non-injective only where required by the boundary conditions. Let $C_1, C_2 \subset B$ be cones over the other two edges of Q with apex at x. The union $Q \cup C_1 \cup C_2$ then patches the hole and joins the 2-cells across γ (note that $\gamma^\phi = x$).

Case $\ell = 1$: We assume that x is an end vertex of γ. Then in the link at x there is a unique point $y := \partial B \cap \gamma^\phi$ corresponding to γ. There also exists a (filled in) triangle $\Delta \subseteq \partial B$ with a vertex mapped to y and its two adjacent edges attached along $\rho_1 \cup \rho_2$, disjoint from the rest of the link, and potentially non-injective only where required by the boundary conditions. Let $C \subset B$ be the cone over the third edge of Δ with its apex at x. The union $\Delta \cup C$ patches the hole and joins the 2-cells across the small initial segment $B \cap \gamma^\phi$ of γ^ϕ. It remains to contract $\gamma^\phi \setminus B$ onto the other end vertex of γ to finalize the joining. □

Joining at paths of length two is no longer guaranteed to preserve embeddability. An example of this failure will be given in the next section once we introduced cloning 2-cells (see Example 3.4).

3.2. **Cloning 2-cells.** Given a 2-cell $c \subseteq X$, *cloning* c means attaching to X a new 2-cell c' with with the same boundary $\partial c' = \partial c$.

Lemma 3.2. *Cloning 2-cells preserves embeddability.*

Proof. Let $\phi\colon X \to \mathbb{R}^4$ be an embedding and $\mathcal{T}$ a triangulation of c^ϕ.

If $\mathcal{T}$ consists of a single triangle Δ then a clone of c can be embedded as follows: let $\sigma \subset \mathbb{R}^4$ be a 3-dimensional simplex with $\sigma \cap X^\phi = \partial\sigma \cap X^\phi = \Delta$. Then $\partial\sigma - \Delta$ is an embedding of the clone c'.

If $\mathcal{T}$ contains more than one triangle then let G be the 1-skeleton of $\mathcal{T}$ and let G' be the subgraph induced on $V(G) \setminus \partial c^\phi$. Let F be a spanning forest of G' and, for each connected component $\tau \subseteq F$, choose an edge $e_\tau \in E(G)$ that connects τ to a boundary vertex $v_\tau \in \partial c^\phi$. Now clone each triangle of $\mathcal{T}$ and join the clones at their shared edges that are not in F. Observe that $\operatorname{int}(c^\phi) \setminus F$ is an open disk and that the previous process created a clone of it. Thus, contracting each connected component $\tau \subseteq F$ onto the corresponding boundary vertex $v_\tau \in \partial c^\phi$ turns these disks into embeddings of the 2-cell c and a clone c'. □

This result has some subtle elements. First, the proof of Lemma 3.2 makes significant use of the PL structure and it is not clear how to translate it to, say, the topological category. Second, to embed the clone our construction modifies the initial 2-cell, and it is not clear whether this can be avoided:

Question 3.3. *Can a clone of a 2-cell be embedded without modifying the embedding of the original 2-cell?*

For a smooth embedding one can always embed a clone in the normal bundle of a 2-cell without modifying the original. In contrast, there are topological embeddings of disks that cannot be completed to embeddings of spheres [2,3], which shows that in the topological category cloning 2-cells will require modification (see also the discussion in [20]). We are not aware of an answer in the PL category.

Cloning is a powerful tool as it allows us to be non-destructive with other operations on 2-cells (such as joining): first clone the involved 2-cells and then perform the operation only on the clones. We will make use of this idea at many instances. For example, below we show that joining 2-cells at a path of length two is not guaranteed to preserve embeddability.

Example 3.4. Starting from $\mathcal{K}_7 - \Delta$ (which is embeddable by Example 2.1), clone the 2-cells $x_1x_2x_4$ and $x_2x_3x_4$ and join the clones at the shared edge x_2x_4. This creates a 2-cell c along $x_1x_2x_3x_4$ and preserves embeddability. However, joining c and a clone of the 2-cell along $x_3x_4x_1$ at the shared 2-path $x_3x_4x_1$ introduced a new 2-cell along $x_1x_2x_3$. We thereby recreated the missing triangle and thus the non-embeddable complex $\mathcal{K}_7$.

3.3. **Application: complexes with few dominating vertices.** In [4] the question was raised whether a 2-complex with a single vertex embeds in $\mathbb{R}^4$. The following result gives a strong affirmative answer with an elementary proof.

Theorem 3.5. *If there are two vertices $v, w \in X$ so that each 2-cell $c \subseteq X$ is incident to at least one of them, then X can be embedded in $\mathbb{R}^4$.*

Proof. Assume first that all 2-cells are incident to the vertex $v \in X$. Fix an embedding ϕ of the 1-skeleton G_X into $\mathbb{R}^3 \times \{0\} \subset \mathbb{R}^4$. Let $C_G \subset \mathbb{R}^3 \times \mathbb{R}_+$ be a cone

over G_X^ϕ, and let $C_u, C_e \subseteq C_G$ denote the subcones over its vertices $u \in G_X$ resp. edges $e \subset G_X$. Each C_u is an edge of C_G, and each C_e is a 2-cell of C_G. To embed a 2-cell $c \subseteq X$ we clone all 2-cells C_e, $e \subset \partial c$ and join the clones at the edges C_u, $u \in \partial c \setminus v$. This results in a single 2-cell $\tilde{c}$ with boundary $\partial\tilde{c} = \partial c \cup C_v$. After having constructed $\tilde{c}$ for all 2-cells $c \subseteq X$, we contract the edge C_v onto v. This results in an embedding of X.

If each 2-cell is incident to one of the two vertices $v, w \in X$, we proceed as above, except that we also embed a second cone C'_G into $\mathbb{R}^3 \times \mathbb{R}_-$. We use C_G to construct the 2-cells that contain v as above, and then use C'_G to construct the remaining 2-cells, which contain w. □

Remark 3.6. If every 2-cell is incident to one of three vertices, then embeddability is no longer guaranteed. The classical counterexamples is the triple-cone over K_5.

3.4. **Rerouting 2-cells.** Let $\gamma_1, \gamma_2 \subseteq G$ be *parallel* paths, that is, with the same end vertices. *Rerouting* a 2-cell c with $\gamma_1 \subset \partial c$ from γ_1 to γ_2 means to replace it by a 2-cell c' that is attached along $(\partial c - \gamma_1) \cup \gamma_2$ instead.

Lemma 3.7. *If γ_1, γ_2 are parallel paths and γ_1 is of length ≤ 1 (i.e., it is a vertex or edge), then the following are equivalent:*

(*i*) *Rerouting a 2-cell from γ_1 to γ_2 yields an embeddable complex.*
(*ii*) *Attaching a 2-cell along $\gamma_1 \cup \gamma_2$ yields an embeddable complex.*

Proof. We show $(i) \implies (ii)$: attaching a new 2-cell $\tilde{c}$ along $\partial\tilde{c} = \gamma_1$ (if γ_1 is an edge then this means $\partial\tilde{c}$ traverses γ_1 twice, in each directions once) preserves embeddability. Since rerouting from γ_1 to γ_2 preserved embeddability, we can reroute a part of $\partial\tilde{c}$ attached along γ_1 to traverse γ_2 instead to create a 2-cell along $\gamma_1 \cup \gamma_2$.

For $(ii) \implies (i)$ let $\tilde{c}$ be a 2-cell attached along $\partial c = \gamma_1 \cup \gamma_2$. To reroute a 2-cell c from γ_1 to γ_2, clone $\tilde{c}$ and join the clone with c at γ_1. Since γ_1 is of length ≤ 1, this preserves embeddability by Lemma 3.1. □

Already rerouting from a path of length two, even with a 2-cell along $\gamma_1 \cup \gamma_2$, is *not* guaranteed to preserve embeddability.

Example 3.8. In $\mathcal{K}_7 - \Delta$ clone the 2-cells $x_1x_2x_4$ and $x_2x_3x_4$ and join the clones at the shared edge x_2x_4. This creates a 2-cell c along $x_1x_2x_3x_4$ and preserves embeddability. The paths $\gamma_1 = x_1x_4x_3$ and $\gamma_2 = x_1x_2x_3$ now satisfy Lemma 3.7 (*ii*) (as c is attached along $\gamma_1 \cup \gamma_2$). However, cloning the 2-cell along $x_1x_3x_4$ and rerouting the clone from γ_1 to γ_2 yields a 2-cell along $x_1x_2x_3$. We thereby recreated the missing triangle and thus the non-embeddable complex $\mathcal{K}_7$.

3.5. **Collapsing 2-cells.** Given a 2-cell $c \subseteq X$, fix a decomposition $\partial c = \gamma_1 \cup \gamma_2$ into parallel paths, as well as a homeomorphism $f\colon \gamma_1 \to \gamma_2$ that fixes end points. *Collapsing* the 2-cell c means deleting it and identifying the paths γ_1 and γ_2 along the homeomorphism f.

Collapsing can be interpreted a an extreme form of rerouting: we reroute every 2-cell from γ_1 to γ_2, even cells that are attached to only a part of γ_1. In contrast to rerouting single 2-cells, where embeddability is preserved only in few special cases, much stronger guarantees can be given when collapsing:

Lemma 3.9. *Collapsing 2-cells preserves embeddability.*

Proof. Let $v_1, ..., v_n \in \gamma_1$ be an enumeration of the vertices of γ_1 and let $e_i \subset \gamma_1$ be the edge between v_i and v_{i+1}. We subdivide c into a "chain of 2-cells" $c_2, ..., c_{n-1} \subseteq c$ by introducing new edges $\tilde{e}_i \subset c$ that connect v_i and $f(v_i)$. Then $\partial c_i = e_i \cup \tilde{e}_i \cup f(e_i) \cup \tilde{e}_{i-1}$. If we contract all $\tilde{e}_i$ then we have $\partial c_i = e_i \cup f(e_i)$. Using Lemma 3.7 $(ii) \implies (i)$ we can reroute all 2-cells attached along e_i to $f(e_i)$. □

For later use we also introduce collapsing cylinders: given a graph H, collapsing a cylinder $H \times [0,1] \subseteq X$ means deleting it and identifying $H \times \{0\}$ with $H \times \{1\}$ in the obvious way.

Corollary 3.10. *Collapsing cylinders preserves embeddability.*

Proof. For each vertex $v \in H$ contract $v \times [0,1]$. This turns $H \times [0,1]$ into a number of 2-cell c_e indexed by edges $e \subseteq H$ and bounded by $\partial c_e = (e \times \{0\}) \cup (e \times \{1\})$. Subsequently collapse each 2-cell c_e onto $e \times \{0\} \subset \partial c_e$. □

3.6. **Cloning and merging edges.** *Cloning* an edge $e \subseteq G$ in a graph means to add in a parallel edge e'. In a complex, cloning e additionally adds the following new 2-cells

- (i) a new 2-cell $\tilde{c}$ along $e \cup e'$, and
- (ii) for each 2-cell $c \subseteq X$ incident to e, another 2-cell c' with the same attachment map as c except that it traverses e' instead of e.

Corollary 3.11. *Cloning edges in complexes preserved embeddability.*

Proof. Fix an embedding $\phi\colon X \to \mathbb{R}^4$. First, embed a disk $D \subset \mathbb{R}^4$ with $D \cap X^\phi = \partial D \cap X^\phi = e^\phi$. We can now interpret $\partial D - e^\phi$ as an embedding of e' and D as an embedding of a 2-cell $\tilde{c}$ with $\partial \tilde{c} = e \cup e'$. For each 2-cell c incident to e we do the following: clone c and reroute the clone from e to e'. This preserves embeddability using Lemmas 3.2 and 3.7. □

We can now confirm van der Holst's Conjecture 1.2 on 4-flat graphs:

Corollary 3.12. *Cloning edges in graphs preserves 4-flatness.*

Proof. Let G' be the graph obtained from G by cloning the edge e. Let X' be the complex obtained from $X(G)$ by cloning the edge e. Observe that $X(G') = X'$. The claim then follows from Corollary 3.11. □

Merging parallel edges $e_1, e_2 \subseteq X$ means to identify them in the obvious way.

Corollary 3.13. *Given parallel edges $e_1, e_2 \subset X$, the following are equivalent:*

- (i) *The complex X' obtained by merging e_1 and e_2 is embeddable.*
- (ii) *The complex X'' obtained by attaching a 2-cell along $e_1 \cup e_2$ is embeddable.*

Proof. Say the new edge in X' is called e. Then the complex obtained from X' by cloning e contains X'' as a subcomplex. Therefore $(i) \implies (ii)$ follows from Corollary 3.11. Conversely, X' is obtained from X'' by collapsing the 2-cell along $e_1 \cup e_2$ onto, say, e_1. Therefore $(ii) \implies (i)$ follows from Lemma 3.9. □

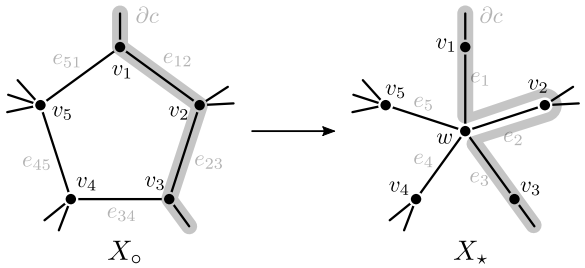

Figure 4. Visualization of stellification of a 5-cycle in a complex. The highlighted paths show how an attachment map ∂c gets modified in the process.

3.7. **Stellifying cycles and ΔY-transformations.** Corollary 3.13 states that, from the point of embeddability, "collapsing a 2-cycle" is the same as filling it by a 2-cell. In this section we investigate whether statements of this sort can be transferred to longer cycles.

Given a cycles $C \subseteq G$ of length $\ell \geq 3$, *stellifying* C mean to replace it by a star (see Figure 4). More precisely, let C have vertices $v_1, ..., v_\ell$ and edges $e_{ij} = v_i v_j$; to stellify C we remove the edges e_{ij} and then add a new vertex w as well as the edges $e_i := wv_i$. If $C \subseteq G_X$ is part of a complex, then stellifying additionally reroutes every 2-cell that traverses e_{ij} to run along $e_i \cup e_j$ instead.

The following notation is useful: we write $G_\circ$ for the graph that contains a cycle C, and $G_\star$ for the graph obtained by stellifying C. Likewise, we write $X_\circ$ and $X_\star$ for the complexes before and after stellifying C. We also write $X_\bullet$ for the complex obtained by attaching a 2-cell along C.

Generalizing Corollary 3.13 to longer cycles turns out non-trivial. We make the following conjecture:

Conjecture 3.14. *If $X_\circ$ is embeddable, then the following are equivalent:*

(*i*) *$X_\star$ is embeddable.*
(*ii*) *$X_\bullet$ is embeddable.*

Lemma 3.15. *Conjecture 3.14 (ii) $\Longrightarrow$ (i) holds.*

Proof. Let $c \subseteq X_\bullet$ be the 2-cell with $\partial c = C$. We subdivide c by adding a vertex w in its interior as well as edges e_i connecting w to v_i. Let $c_i \subset c$ be the sub-cell with $\partial c_i = e_i \cup e_{ij} \cup e_j$. Stellifying C is then achieved by collapsing each c_i onto $e_i \cup e_j$. This preserves embeddability by Lemma 3.9. □

So far we are unable to prove the other direction (*i*) $\Longrightarrow$ (*ii*). We also emphasize an important difference to the 2-cycle version (Corollary 3.13): in Conjecture 3.14 we need to assume the embeddability of $X_\circ$, as otherwise the direction (*i*) $\Longrightarrow$ (*ii*) is not true. An example of this failure is given in Section 4.1.

A particularly common instance of stellification is $\ell = 3$, in which case the operation is known as a ΔY-*transformation*. Using this terminology Lemma 3.15 reads

Corollary 3.16. *ΔY-transformations in complexes preserve embeddability.*

We next work towards van der Holst's Conjecture 1.3 for 4-flat graphs. Let us use the notation G_Δ and G_Y instead of $G_\circ$ and $G_\star$ to emphasize that we work with $\ell = 3$ (and X_Δ and X_Y accordingly). We make the following observation:

Observation 3.17. If $X_\Delta := X(G_\Delta)$, then $X_\mathrm{Y} = X(G_\mathrm{Y})$. This statement is slightly more subtle than one might recognize at first. In fact, the analogous statement is incorrect for cycles of lengths $\ell \geq 4$: while in $X(G_\star)$ there is a 2-cell that traverses $e_1 \cup e_3$, there is no such 2-cell in $X_\star$.

The ΔY-part of van der Holst's Conjecture 1.3 follows immediately:

Corollary 3.18. *ΔY-transformations in graphs preserve 4-flatness.*

Proof. Set $X_\Delta := X(G_\Delta)$. As observed in Observation 3.17 we have $X_\mathrm{Y} = X(G_\mathrm{Y})$. The claim then follows from Corollary 3.16. □

Observation 3.19.

(*i*) By cloning all edges of the cycle C we create a second cycle C' on the same vertex set. Stellifying this new cycle has the net effect of adding to G a suspension vertex over C while also preserving the original cycle.

(*ii*) Since both cloning edges and ΔY-transformations preserve 4-flatness, from (*i*) follows that adding a suspension vertex over a triangle in G preserves 4-flatness. We prove a strengthening of this in Lemma 5.14.

(*iii*) Adding a cone over a cycle and then deleting all but two of the cone edges has the net effect of adding a chord to the cycle. Any chord can be created in this way.

(*iv*) We can now see that stellifying at cycles of length $\ell \geq 4$ cannot preserve 4-flatness: via cloning edges and stellifying at 4-cycles one can reconstruct the missing edge in the 4-flat graph $K_7 - e$, thereby turning it into K_7, which is not 4-flat.

Consider the following generalized stellification operation: stellifying a subgraph $H \subseteq G$ means to delete its edges and to add a new suspension vertex over it (and note that with the help of doubling edges, deleting the edges of H does not make any difference for 4-flatness). Following the reasoning of Observation 3.19 (*iii*) and (*iv*), stellifying a non-complete subgraph cannot preserve 4-flatness. Likewise, stellifying K_5 cannot preserve 4-flatness either, because we can use it to turn K_5 into a triple suspension over K_5, which is not 4-flat. The following question remains open:

Question 3.20. *Does stellifying K_4 preserve 4-flatness?*

4. Reverse stellification and YΔ-transformations

In Section 3.7 we explored the conditions under which the embeddability of $X_\circ$ implies the embeddability of $X_\star$, as well as analogous questions for 4-flatness. In this section we ask about the opposite direction: if $X_\star$ embeds, what can we say about $X_\circ$? This includes the YΔ-part of van der Holst's Conjecture 1.3:

Conjecture 4.1. *YΔ-transformations in graphs preserve 4-flatness.*

As of now, this conjecture remains open. We do however present evidence that its statement is right at the boundary of what can be true:

(*i*) an analogue statement is not true for cycles of length $\ell \geq 4$: even if $G_\star$ is 4-flat, $G_\circ$ might not (see Example 4.2 below).

(*ii*) an analogue statement is not true for embeddability of general complexes: even if X_Y is embeddable, X_Δ might not (see Section 4.1).

Example 4.2. $G_\circ := K_7$ is not 4-flat. However, stellifying any of its 4-cycles results in a 4-flat graph $G_\star$. To see this, we make use of the fact that the double suspension of a planar graph is 4-flat (we prove this in Theorem 5.9 in the next section). Since we can delete two vertices of $G_\star$ and obtain a planar graph (see Figure 5), $G_\star$ is contained in such a double suspension. In conclusion, reverse stellification at a 4-cycle does not always preserve 4-flatness.

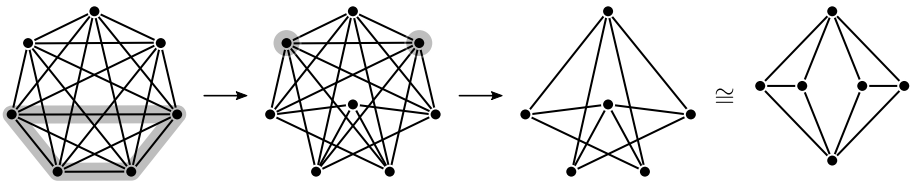

Figure 5. Starting from K_7 we first stellify a 4-cycle, then we delete two vertices to obtain a planar graph.

4.1. **A YΔ-transformation that does not preserve embeddability.** The following construction is based on the Freedman-Krushkal-Teichner (FKT) complex X_{FKT} [6]. One of the main characteristics of the FKT complex is that it does *not* embed into $\mathbb{R}^4$, despite its van Kampen obstruction being zero.

We shall describe two complexes, X_Δ and X_{Y}, the latter being the ΔY-transformation of the former. We will show that X_{Y} is embeddable, while X_Δ is not (by proving that any potential embedding of it can be turned into an embedding of the FKT complex).

We briefly recall the FKT complex and then modify it to obtain X_Δ.

4.1.1. *The Freedman-Krushkal-Teichner complex.* The complex is built as follows:

(1) start from two copies of K_7, say, K_7 with vertices $x_1, ..., x_7$, and $\tilde{K}_7$ with vertices $\tilde{x}_1, ..., \tilde{x}_7$.
(2) attach a 2-cell to each 3-cycle of K_7, except $x_4x_5x_6$. Analogously, attach a 2-cell to each 3-cycle of $\tilde{K}_7$, except $\tilde{x}_4\tilde{x}_5\tilde{x}_6$.
(3) add an edge e_6 between x_6 and $\tilde{x}_6$.
(4) attach a 2-cell c^* along the closed walk ∂c^* given by

$$x_6\, x_5\, x_4\, x_6\, \tilde{x}_6\, \tilde{x}_5\, \tilde{x}_4\, \tilde{x}_6\, x_6\, x_4\, x_5\, x_6\, \tilde{x}_6\, \tilde{x}_4\, \tilde{x}_5\, \tilde{x}_6. \tag{4.1}$$

The contrived attachment map ∂c^* is chosen specifically to force the emergence of a Borromean ring like structure in every potential embedding of the complex (details can be found in the original work of Freedman, Krushkal and Teichner [6], or with more explanations in [1]). The relevant property of X_{FKT} for our purpose is the following (see [6, Section 3.2]): no embedding of $X_{\mathrm{FKT}} - c^*$ can be extended to a mapping of X_{FKT} into $\mathbb{R}^4$, which means that we are allowing self-intersections c^*, but still no intersections between c^* and other parts of X_{FKT}.

4.1.2. *The complexes X_Δ and X_{Y}.* We modify X_{FKT} to obtain the complex X_Δ:

(1) identify c_{123} and $\tilde{c}_{123}$ in an interior point; call the resulting point w.
(2) identify c^* and c_{123} in an interior point; call the resulting point v.
(3) identify c^* and $\tilde{c}_{123}$ in an interior point; call the resulting point $\tilde{v}$.

The three identifications do not respect the CW-complex structure, which needs to be restored using suitable subdivisions. We choose subdivisions which in particular contain the following edges:

(4) an edge $e_{123} \subset \operatorname{int}(c_{123})$ from v to w.
(5) an edge $\tilde{e}_{123} \subset \operatorname{int}(\tilde{c}_{123})$ from $\tilde{v}$ to w.
(6) an edge $e^* \subset \operatorname{int}(c^*)$ from v to $\tilde{v}$.

The edges $e_{123}, \tilde{e}_{123}$ and e^* form the triangle Δ on which we later perform the ΔY-transformation. The complex obtained by this ΔY-transformation is X_{Y}.

4.1.3. *Non-embeddability of X_Δ.* In [6, Section 3.2] the authors prove that an embedding of $X_{\mathrm{FKT}} - c^*$ cannot be extended to a mapping of X_{FKT} into $\mathbb{R}^4$, in particular, allowing for self-intersections of c^* (a fact that we will make use of). By construction, an embedding of X_Δ defines a mapping of X_{FKT} into $\mathbb{R}^4$ in which c_{123}, $\tilde{c}_{123}$ and c^* have pairwise intersections at $v, \tilde{v}$ and w respectively, but that is an embedding otherwise. We show that we can get rid of these intersections, creating self-intersections at most for c^*, and in this way, turning an embedding of X_Δ into a mapping of X_{FKT} that is known to not exist.

Lemma 4.3. *Let $\phi_1, \phi_2 \colon \mathbb{S}^2 \to \mathbb{R}^4$ be* PL *maps, potentially non-injective, with a single intersection $x := \phi_1 \cap \phi_2$, and injective in a neighborhood B of x. Then the intersection can be removed by modifying the ϕ_i in this neighborhood of x. If ϕ_i was injective to begin with, then it is still injective after the modification.*

Proof. Delete $(\phi_1 \cup \phi_2) \cap B$ and replace it by a copy of the exterior $(\phi_1 \cup \phi_2) \setminus B$ via inversion on the 3-sphere ∂B. This replaces the disk $\phi_i \cap B$ by an image of the disk $\phi_i \setminus B$, keeping ϕ_i a map of the 2-sphere, and injective if ϕ_i was injective outside of B. Since there are no intersections of ϕ_1 and ϕ_2 outside of B, this removed the intersection. □

Observe that the sub-complex σ induced on the vertices $x_1, x_2, x_3, x_7 \in X_\Delta$ has 2-cells along each triangle (like a 3-simplex) and forms a 2-sphere. Analogously the sub-complex $\tilde{\sigma}$ on $\tilde{x}_1, \tilde{x}_2, \tilde{x}_3, \tilde{x}_7$ forms a 2-sphere. In X_Δ these 2-spheres intersect exactly once in $w = c_{123} \cap \tilde{c}_{123}$. In the embedding this intersection can then be removed using Lemma 4.3.

Next, we observe that the boundary curve ∂c^* can be filled in by a (self-intersecting) disk disjoint from σ: use $\tilde{\sigma} - \tilde{c}_{123}$ to fill in the parts $\tilde{x}_6\tilde{x}_5\tilde{x}_4$ and $\tilde{x}_6\tilde{x}_4\tilde{x}_5$; the rest can be filled in by a "collapsed disk" (see Figure 6). Together with the embedding of

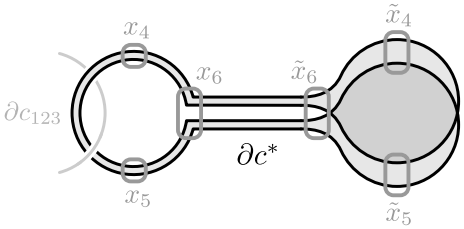

Figure 6. Visualization of a disk attached along ∂c^* as constructed in Section 4.1.3. Note that ∂c^* is non-injective but is depicted in an injective way to assist the visualization. The curve passes through a single vertex in the regions marked in gray.

c^* this is a mapping (but not an embedding) of a 2-sphere. We now use Lemma 4.3 to get rid of the unique intersection of c^* and c_{123} at v, potentially creating self-intersections of c^*. Analogously, we can get rid of the single intersection between c^* and $\tilde{c}_{123}$ at $\tilde{v}$.

In this way we obtained a mapping of X_{FKT} into $\mathbb{R}^4$, having self-intersections only of c^* and being an embedding otherwise. This is a contradiction.

4.1.4. *Embeddability of* X_{Y}. Recall that the complex $\mathcal{K}_7$ is obtained from the complete graph K_7 by attaching a 2-cell along each triangle. Let $x_1, ..., x_7$ be its vertices and let c_{ijk} denote the 2-cell attached along $x_ix_jx_k$. Recall further that while $\mathcal{K}_7$ does not embed, there exists a mapping $\phi\colon \mathcal{K}_7 \to \mathbb{R}^4$ that is injective except for a single intersection between, say, c_{123} and c_{456}. Define the complex $\mathcal{K}_7^*$ from $\mathcal{K}_7$ by identifying c_{123} and c_{456} in an interior point. We denote the point of intersection by $y := c_{123} \cap c_{456}$ and subdivide the intersecting 2-cells by edges of the form yx_i to restore the CW complex structure. We continue to denote the subdivided 2-cells by c_{123} and c_{456} respectively. We can now view ϕ as an embedding of $\mathcal{K}_7^*$. Observe that the link at y^ϕ consists of two linked cycles. We now modify $\mathcal{K}_7^*$ into X_{Y} by a series of operations that preserve embeddability (*cf.* Section 3).

Choose a 4-ball $B \subset \mathbb{R}^4$ that intersects $(\mathcal{K}_7^*)^\phi$ only in y. By inversion on ∂B we embed a copy $\tilde{\mathcal{K}}_7^*$ that shares with $\mathcal{K}_7^*$ only the vertex y. We denote its vertices by $\tilde{x}_1, ..., \tilde{x}_7$ and its 2-cells by $\tilde{c}_{ijk}$ respectively. Observe that the link at y^ϕ now consists of four cycles that belong to two linked pairs but are otherwise unlinked.

Embed a disk $D \subset \mathbb{R}^4$ with a path in its boundary ∂D attached along x_6yx_6. The opposite path $\partial D \setminus (x_6y\tilde{x}_6)^\phi$ we now consider as an embedding of an edge $e_6 = x_6\tilde{x}_6$ and D as a 2-cell attached along $yx_6 \cup y\tilde{x}_6 \cup e_6$. After these additions, the link at y^ϕ now looks as shown in Figure 7.

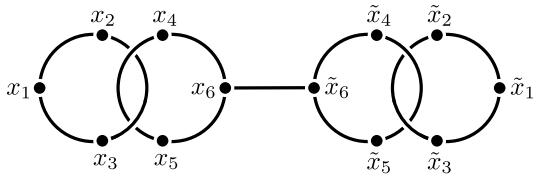

Figure 7. The link at y^ϕ at a stage during the construction of X_{Y}. The point labelled x_i (resp. $\tilde{x}_i$) is the vertex of the link that corresponds to the edge yx_i (resp. $y\tilde{x}_i$) in the complex.

By cloning relevant 2-cells incident to y (*cf.* Lemma 3.2) and joining the clones suitably at shared edges (*cf.* Lemma 3.1) we create an embedding of a (subdivided) 2-cell attached along the path ∂c^* (see also (4.1)) that intersects c_{123} and $\tilde{c}_{123}$ only in y. Figure 8 visualizes the modification to the link at y^ϕ.

It remains to turn y into a vertex of degree three (the center of the Y in X_{Y}). As before, we attach suitable disks $D_1, D_2, D_3 \subset \mathbb{R}^4$ to the complex to modify the link as shown in Figure 9 (top). We then collapse D_i onto edges that we denote as yz_i (see Figure 9 bottom). The z_i form the neighbors of y in X_{Y}. We eventually adjust the subdivisions of the 2-cells so that y is truly of degree three. This finalizes the construction of X_{Y} and its embedding.

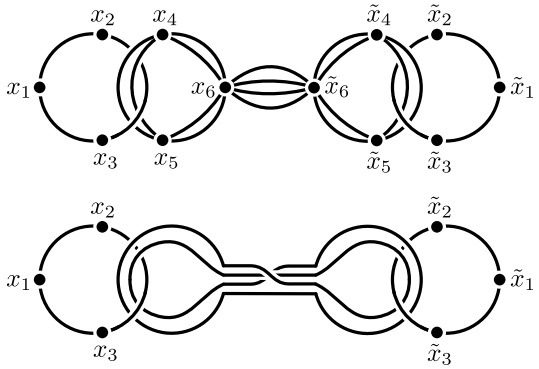

Figure 8. The link of y^ϕ after cloning (top) and then joining (bottom) certain 2-cells. The three cycles form Borromean rings.

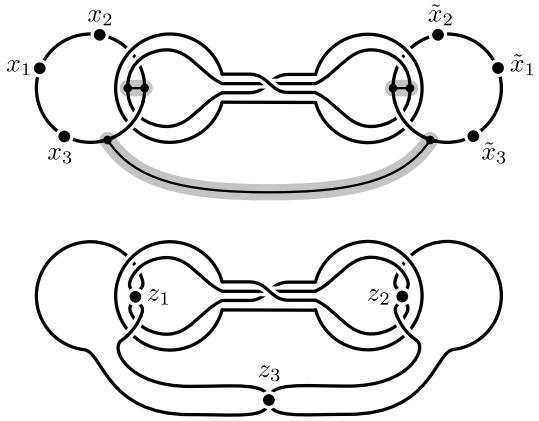

Figure 9. Modification to the link of y^ϕ during the last step of the construction. Link vertices of degree two can be removed by modifying the subdivision of 2-cells adjacent to y.

5. Results on 4-flat graphs

In Section 3 we developed sufficient machinery to already verify van der Holst's Conjecture 1.2 as well as the ΔY-part of Conjecture 1.3. In this section we explore further consequences of our results for the theory of 4-flat graphs.

5.1. **Variants of full complexes.** Recall that the *full complex* $X(G)$ of a graph G has 1-skeleton G and a 2-cell attached along each cycle of G. Consider the following variants:

- (i) $X_{\text{ind}}(G)$ has a 2-cell attached along each *induced* cycle of G.
- (ii) $X_{\text{reg}}(G)$ has a 2-cell attached along each cycle of G (this is just $X(G)$).
- (iii) $X_{\text{full}}(G)$ has a 2-cell attached along each *closed walk* in G.

The notion of 4-flatness discussed so far is defined by the embeddability of $X_{\text{reg}}(G)$. It is natural to consider analogous notions of 4-flatness defined using $X_{\text{ind}}(G)$ and $X_{\text{full}}(G)$ respectively. We show that they are equivalent:

Theorem 5.1. *The following are equivalent:*

(*i*) $X_{\mathrm{ind}}(G)$ *embeds.*
(*ii*) *every regular complex on* G *embeds (or equivalently,* $X_{\mathrm{reg}}(G)$ *embeds).*
(*iii*) *every finite (potentially non-regular) complex on* G *embeds (or equivalently, every finite subcomplex of* $X_{\mathrm{full}}(G)$ *embeds).*

Proof. $(iii) \Longrightarrow (ii) \Longrightarrow (ii)$ is clear. It suffices to verify the other directions.

Every cycle $C \subseteq G$ can be written as a cycle sum $C_1 \oplus \cdots \oplus C_r$ of induced cycles. A 2-cell attached along C can be constructed by cloning 2-cells attached along the C_i and joining the clones at shared edges. Since cloning and joining at edges preserves embeddability (*cf.* Lemma 3.1), we proved $(i) \Longrightarrow (ii)$.

A 2-cell with constant attachment map as well as a 2-cell attached along a single edge (traversing it twice in opposite directions) can always be added to a complex while preserving embeddability. Any other walk γ in G can be written as the concatenation of cycles and double-traversals of edges; call them $\gamma_1, ..., \gamma_r$. A 2-cell attached along γ can be constructed by cloning 2-cells attached along the γ_i and then joining the clones at the concatenation vertices. Since cloning and joining at vertices preserves embeddability (*cf.* Lemma 3.1), we proved $(ii) \Longrightarrow (iii)$. □

Since $X_{\mathrm{ind}}(G)$, $X_{\mathrm{reg}}(G)$ and $X_{\mathrm{full}}(G)$ are equivalent from the point of embeddability, we will continue to write $X(G)$ to denote any of them and choose the interpretation most convenient for the situation. For example, $X_{\mathrm{ind}}(G)$ is most convenient to establish embeddability since it has the fewest 2-cells:

Example 5.2. Since all induced cycles of $K_7 - e$ are triangles we find $X_{\mathrm{ind}}(K_7 - e)$ as a subcomplex of $\mathcal{K}_7 - \Delta$. Since $\mathcal{K}_7 - \Delta$ is embeddable (see Example 2.1), this gives a short proof that $K_7 - e$ is indeed 4-flat.

We present another argument based on Schlegel diagrams of polytopes:

Example 5.3. K_6 is the 1-skeleton of the 5-dimensional simplex. Its 2-dimensional faces form exactly the 2-cells of $X_{\mathrm{ind}}(K_6)$. The Schlegel diagram of the simplex is a 3-dimensional complex embedded in $\mathbb{R}^4$ that contains $X_{\mathrm{ind}}(K_6)$ as a subcomplex. We conclude that K_6 is 4-flat.

Analogously, consider gluing two 5-dimensional simplices at a 4-dimensional face. The resulting polytope has $K_7 - e$ as its 1-skeleton (the missing edge connected the vertices opposite to the glued faces) and each 2-dimensional face corresponds to a 2-cell in $X_{\mathrm{ind}}(K_7 - e)$. Following the argument above, we find $K_7 - e$ is 4-flat.

5.2. **Sliced embeddings.** We say that an embedding of a complex X is *sliced* if it embeds the 1-skeleton G_X into the 3-dimensional subspace $\Pi := \mathbb{R}^3 \times \{0\} \subset \mathbb{R}^4$. A number of 4-flat variants can be defined based on such embeddings. First, we note the following:

Lemma 5.4. *Every embeddable complex has a sliced embedding.*

Proof. For a generic embedding $\phi\colon X \to \mathbb{R}^4$ the projection $\pi\colon \mathbb{R}^4 \to \Pi$ is injective on the embedded skeleton G_X^ϕ. Let $H := \pi(G_X^\phi) \subset \Pi$ be the image of the projection. Choose a triangulation $\mathcal{T}$ of Π that contains H as a subcomplex. For each simplex $\sigma \in \mathcal{T}$ let $\hat{\sigma} := \pi^{-1}(\sigma) = \sigma \times \mathbb{R}$ be the cylinder over the simplex. We define a homeomorphism $f\colon \mathbb{R}^4 \to \mathbb{R}^4$ as follows: for a vertex $v \in \mathcal{T}$, either $v \notin H$ and f fixes the ray $\hat{v}$, or $v \in H$ and f translates the ray $\hat{v}$ by moving $\pi^{-1}(v)$ to v. Then extend f linearly to all cylinders $\hat{\sigma}$. The image $f(X^\phi)$ is a sliced embedding of X. □

This suggests to consider the following subtypes of sliced embeddings:

(*i*) a *+-sliced embedding* embeds each 2-cell into the halfspace $\mathbb{R}^3 \times \mathbb{R}_+$.
(*ii*) a *±-sliced embedding* embeds each 2-cell into one of the halfspaces $\mathbb{R}^3 \times \mathbb{R}_+$ and $\mathbb{R}^3 \times \mathbb{R}_-$.

Both types of sliced embeddings are likely much nicer to work with than general (sliced) embeddings. It turns out that graphs for which $X(G)$ has a +-sliced embedding are exactly the linkless graphs: recall that a graph G is *linkless* if there is an embedding $\phi\colon G \to \mathbb{R}^3$ so that any two disjoint cycles of G are mapped to unlinked closed curves.

Theorem 5.5. *The following are equivalent:*

(*i*) *$X(G)$ has a +-sliced embedding.*
(*ii*) *G is linkless.*

Proof. The proof of $(ii) \implies (i)$ presented below is essentially due to van der Holst [17, Theorem 2] but we include it here for completeness and for its elegance. If G is linkless, then it has a flat embedding $\phi\colon G \to \mathbb{R}^3$ [13]. That means, for each 2-cell $c \subseteq X_{\text{reg}}(G)$ (which is attached along a cycle in G) there exists an embedded disk $D_c\colon c \to \mathbb{R}^3$ with $D_c \cap G^\phi = \partial D_c \cap G^\phi = \partial c$. We choose distinct numbers $a_c > 0$, one for each 2-cell $c \subseteq X_{\text{reg}}(G)$. We now extend ϕ to embed the 2-cell c into $\mathbb{R}^3 \times \mathbb{R}_+$ by the following map:

$$c \ni x \;\mapsto\; \phi(x) := \begin{pmatrix} D_c(x) \\ a_c \operatorname{dist}(D_c(x), G^\phi) \end{pmatrix} \in \mathbb{R}^3 \times \mathbb{R}_+$$

where $\operatorname{dist}(D_c(x), G^\phi)$ is the distance of $D_c(x)$ to the closest point in G^ϕ. It remains to show that this is an embedding. Clearly, ϕ is an embedding on G with any single 2-cell. If two embedded 2-cells c_1^ϕ and c_2^ϕ were to intersect in points $\phi(x_1) = \phi(x_2)$, where $x_i \in c_i$, then comparing the $\phi(x_i)$ component-wise yields

$$a_{c_1} \operatorname{dist}(D_{c_1}(x_1), G^\phi) = a_{c_2} \operatorname{dist}(D_{c_2}(x_2), G^\phi) = a_{c_2} \operatorname{dist}(D_{c_1}(x_1), G^\phi),$$

and hence $a_{c_1} = a_{c_2} \implies c_1 = c_2$, which is a contradiction.

The direction $\neg(ii) \implies \neg(i)$ follows from two facts. First, if G is not linkless, then each embedding in $\mathbb{R}^3$ contains two disjoint cycles of non-zero linking number (by definition it contains two linked cycles, but the non-zero linking number follows from the excluded-minor characterization of linkless graphs and the fact that such a pair of cycles exist for each excluded minor; see [15] or [19, Theorem 36]). Second, a pair of disjoint closed curves in $\mathbb{R}^3 \times \{0\}$ of non-zero linking number cannot be filled in by disjoint disks in $\mathbb{R}^3 \times \mathbb{R}_+$ (see *e.g.* [16, Lemma 2]). □

In contrast, graphs G for which $X(G)$ has a ±-sliced embeddings appear to be a much more comprehensive class. While we suspect that not every embeddable complex (or full complex of a 4-flat graph) has a ±-sliced embedding, so far we do not know any examples.

Question 5.6.

(*i*) *Has every embeddable complex a ±-sliced embedding?*
(*ii*) *If G is 4-flat, has its full complex $X(G)$ a ±-sliced embeddings?*

5.3. **4-flat and linkless graphs.** We say that a graph is *locally linkless* if the neighborhood $N_G(x)$ of each vertex $x \in G$ is a linkless graph.

Lemma 5.7. *4-flat graphs are locally linkless.*

Proof. Fix an embedding $\phi\colon X(G) \to \mathbb{R}^4$. For a vertex $x \in G$, let $B \subset \mathbb{R}^4$ be a link neighborhood of x^ϕ. Note that $X(G)$ contains a cone C over $N_G(x)$ with apex at x. Then $\mathrm{cl}(C^\phi \setminus B)$ is a cylinder of the form $N_G(x) \times [0,1]$. Collapsing the cylinder onto $C \cap \partial B$ preserves embeddability by Corollary 3.10. This results in an embedding of $X(G)$ for which $N_G(x)$ is embedded into ∂B, and each 2-cell non-incident to x is embedded outside of B. In particular, it yields a $+$-sliced embedding of $X(N_G(x))$. $N_G(x)$ must then be linkless by Theorem 5.5. □

We obtain a quick argument for the following well-known fact:

Corollary 5.8. *K_7 and $K_{3,3,1,1}$ are not 4-flat.*

Proof. The neighborhood of a vertex in K_7 is K_6. The neighborhood of a dominating vertex in $K_{3,3,1,1}$ is $K_{3,3,1}$. The graphs K_6 and $K_{3,3,1}$ are not linkless. □

5.4. **4-flatness and suspensions.** Given graphs G and H, by $G * H$ we denote the graph obtained from the disjoint union $G \cup H$ by adding a complete bipartite graph between the vertices of G and the vertices of H. For example, $G * K_1$ is a single suspension of G, and $G * K_n$ is an n-fold (iterated) suspension.

Theorem 5.9.

(*i*) *$G * K_1$ is 4-flat if and only if G is linkless.*
(*ii*) *$G * K_2$ is 4-flat if and only if G is planar.*
(*iii*) *$G * K_3$ is 4-flat if and only if G is outerplanar.*

Proof. A graph is outerplanar if and only if its suspension is planar; and a graph is planar if and only if its suspension is linkless [18]. It therefore suffices to prove (*i*).

Let x be the suspension vertex of $G * K_1$. If $G * K_1$ is 4-flat, then linklessness of G follows from $G \simeq N_{G*K_1}(x)$ and Lemma 5.7. For the converse, observe that $X_{\mathrm{ind}}(G * K_1)$ is obtained from $X_{\mathrm{ind}}(G)$ by adding a cone over G. If G is linkless, then by Theorem 5.5 there exists a $+$-sliced embedding $\phi\colon X_{\mathrm{ind}}(G) \to \mathbb{R}^3 \times \mathbb{R}_+$. We can extend ϕ to an embedding of $X_{\mathrm{ind}}(G * K_1)$ by embedding the a cone over G into $\mathbb{R}^3 \times \mathbb{R}_-$. Thus, $G * K_1$ is 4-flat. □

The join operation can be used to define potentially interesting intermediate classes between outerplanar, planar and linkless graphs.

Example 5.10. $G * \bar{K}_2$ is 4-flat for some linkless graphs (*e.g.* $G = K_5$ gives $G * \bar{K}_2 = K_7 - e$), but not for others (*e.g.* $G = K_{3,1,1,1}$ gives $G * \bar{K}_2 = K_{3,3,1,1}$). Likewise, $G * \bar{K}_3$ is 4-flat for some planar graphs (*e.g.* $G = K_4$ gives $G * \bar{K}_3 = K_7 - \Delta$), but not for others (*e.g.* $G = K_{3,1,1}$ gives $G * \bar{K}_3 = K_{3,3,1,1}$).

Problem 5.11. *Given H, describe the class of graphs G for which $G * H$ is 4-flat.*

5.5. **4-flat and knotless graphs.** A graph G is *knotless* if it has a spacial embedding $\phi\colon G \to \mathbb{R}^3$ for which each cycle $C \subseteq G$ is embedded as the trivial knot. A graph that has no such embedding is called *intrinsically knotted*. The knotless graphs are often considered as another natural "higher analogue" of linkless graphs and one might ask about their relation to 4-flat graphs. The following two examples demonstrate that these two classes are generally unrelated:

Example 5.12. While the ΔY-family of $K_{3,3,1,1}$ consists entirely of intrinsically knotted graphs, seven out of the 20 members of the ΔY-family of K_7 are in fact linkless [7, Table 1]. Yet, as Heawood graphs, none of them is 4-flat. Thus, there are graphs that are knotless while not being 4-flat.

Example 5.13. There exists a linkless graph G whose suspension $G * K_1$ is not knotless [5]. However, by Theorem 5.9 (i), $G * K_1$ is 4-flat. Thus, there are graphs that are 4-flat while not being knotless.

5.6. **Clique sums of 4-flat graphs.** Given two graphs G_1 and G_2, both of which contain a k-clique K_k, by the *k-clique sum* $G_1 *_k G_2$ we mean the graph obtained by "gluing" the graphs at their designated K_k subgraphs.

Lemma 5.14. *4-flat graphs are closed under 3-clique sums.*

Proof. Let G_1 and G_2 be two 4-flat graphs with 3-cliques $\Delta_i \subseteq G_i$ and embeddings $\phi_i : X(G_i) \to \mathbb{S}^4$. Let $c_i \subseteq X(G_i)$ be the 2-cell with $\partial c_i = \Delta_i$ and $x_i \in \mathrm{int}(c_i)$ an interior point at which the embedding is locally flat. That is, given a link neighborhood $B_i \subset \mathbb{R}^4$ of $x_i^{\phi_i}$, the link $\partial B_i \cap c_i^{\phi_i}$ is unknotted in ∂B_i. There is then a homeomorphism $(\partial B_1, \partial B_1 \cap c_1^{\phi_1}) \simeq (\partial B_2, \partial B_2 \cap c_2^{\phi_2})$. We identify $\mathbb{S}^4 \setminus B_1$ and $\mathbb{S}^4 \setminus B_2$ and their embedded complexes $(X(G_i), \phi_i)$ along this homeomorphism, which yields a new $\mathbb{S}^4$ with an embedded complex X'. This new complex X' contains a cylinder with boundary $\Delta_1 \cup \Delta_2$. We collapse the cylinder (*i.e.,* we identify Δ_1 and Δ_2), which preserves embeddability by Corollary 3.10. The resulting complex X'' has 1-skeleton $G_1 *_3 G_2$. Since induced cycles in $G_1 *_3 G_2$ do not pass through the shared 3-clique, we find $X_{\mathrm{ind}}(G_1 *_3 G_2) \subseteq X''$. $G_1 *_3 G_2$ is then 4-flat by Theorem 5.1. □

Question 5.15. *Do 4-clique sums preserve 4-flatness?*

This question generalizes Question 3.20.

Note that 5-clique sums do *not* preserve 4-flatness: starting from three disjoint copies of K_6 (which is 4-flat) we can use 5-clique sums to glue them at K_5-subgraphs to obtain $K_5 * \bar{K}_3$. The full complex of this graph contains the triple cone over K_5, which is not embeddable.

6. Heawood graphs are excluded minors

Van der Holst conjectured that the Heawood graphs are exactly the excluded minors for the class of 4-flat graphs (Conjecture 1.1). Proving this statement appears extraordinarily hard and especially proving that there are no other excluded minors is likely out of reach of current techniques. However, even though van der Holst showed that the Heawood graphs are not 4-flat, so far it has also not been verified that they are excluded minors, *i.e.,* that all their minors are 4-flat. The main reason for this is that this claim appears not accessible by a concise or systematic argument, but rather comes down to a lengthy case analysis, checking each minor of each of the 78 Heawood graphs for 4-flatness.

In this section we verify this claim in two ways and establish

Theorem 6.1. *All graphs of the Heawood family are excluded minors for the class of 4-flat graphs.*

Both approaches are based on case analysis. The first proof relies on computer help, actually enumerating all minors of all Heawood graphs and checking them one

by one (see Section 6.1). The second proof uses results from Section 5 to reduce the case analysis to merely five Heawood graphs, which can then be checked by hand (see Section 6.2).

The general procedure for both proofs is as follows: given a Heawood graph G, we step through its minors H (where it suffices to consider minors of the form $G-e$ and G/e) and show that each one is 4-flat. We do so by identifying two vertices $v, w \in V(H)$ for which $H' := H - \{v, w\}$ is planar. This shows that H is a subgraph of $H' * K_2$ and therefore 4-flat by Theorem 5.9 (*ii*).

Note that this approach for detecting 4-flat graphs is not guaranteed to succeed a priorily: there are 4-flat graphs that are not contained in the double suspension of a planar graph (*e.g.* the disjoint union of two K_6). Computationally we found that this works at least for the minors of Heawood graphs. The following question remains:

Question 6.2. *If G is an excluded minor for the class of 4-flat graphs and H is a minor of G, then are there two vertices $v, w \in V(H)$ so that $H - \{v, w\}$ is planar?*

6.1. **Proof by computer.** Code for both enumerating all Heawood graphs and performing the exhaustive case analysis can be found in Appendix A. The output of the program confirms Theorem 6.1.

6.2. **Proof by hand.** In Section 3 we proved that 4-flat graphs are closed under cloning edges (Corollary 3.12) and ΔY-transformations (Corollary 3.18). This constitutes an affirmative answer to van der Holst's conjecture Conjecture 1.2 and part of Conjecture 1.3. Assuming the truth of both conjectures, van der Holst showed that all Heawood graphs are indeed excluded minors for the class of 4-flat graphs [17, Lemma 1]. However, since we only verified the ΔY-part of Conjecture 1.3 we cannot draw this conclusion right away. Our results do however sufficiently reduce the work that remains to be done, so as to allow for a proof by hand.

The following lemma is a key tool used to reduce the case analysis to only a hand full of graphs:

Lemma 6.3. *Let G be an excluded minor for the class of 4-flat graphs and let H be a graph obtained from G by a ΔY-transformation. Then either H is 4-flat, or H is itself an excluded minor for the class of 4-flat graphs.*

Proof. It suffices to show that either H is 4-flat, or every proper minor H' of H is 4-flat. Since the class of 4-flat graphs is minor-closed, it suffices to consider the cases $H' = H - e$ and $H' = H/e$ for some edge $e \in E(H)$.

Let $x, y, z \in V(H)$ be the vertices of the triangle $\Delta \subset G$ on which we performed the ΔY-transformation, and let v by the resulting vertex of degree three in H. There are two cases to be considered.

Case 1: e does *not* contain v. Then e is also an edge in G, Δ is also a triangle in $G' := G - e$ resp. $G' := G/e$, and H' is obtained from G' via ΔY-transformation on Δ. Since G is an excluded minor, G' is 4-flat. H' is then 4-flat by Corollary 3.18.

Case 2: e does contain v. We assume w.l.o.g. that $e = vx$. Then H/e coincides with $G - yz$, which is 4-flat. Similarly, $H - e$ is obtained from $G - \{xy, xz\}$ by subdividing the edge yz, which again is 4-flat. □

By Lemma 6.3 it suffices to consider the Heawood graphs that are not ΔY-transformations of other Heawood graphs. If those are shown to be excluded minors for the class of 4-flat graphs, then it already follows that all Heawood graphs are.

Note that these are exactly the Heawood graphs of minimum degree ≥ 4 (because we cannot perform a YΔ-transformation on them). Code for listing these Heawood graphs can be found in Appendix A.3. This list contains K_7 and $K_{3,3,1,1}$ (which are known to be excluded minors; [17, Lemma 2]) and only five other graphs (shown in Figure 10). We call them the *remaining Heawood graphs*.

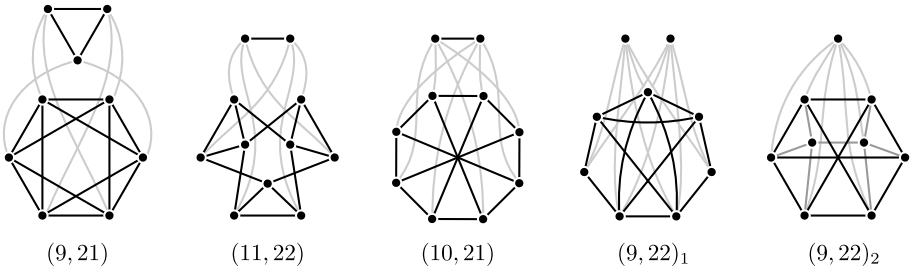

Figure 10. The five remaining Heawood graphs. The edges are colored solely for visualization purpose. The name (v, e) indicates that the graph as v vertices and e edges, with a sub-index to resolve disambiguities.

In the subsequent subsections we treat the remaining Heawood graphs one-by-one:

6.2.1. *The case* $(9, 21)$.

There are at most four orbits of edges as coloured in Figure 11 (a). The bottom half is the square of a hexagon, which has a planar embedding H as shown in Figure 11 (b).

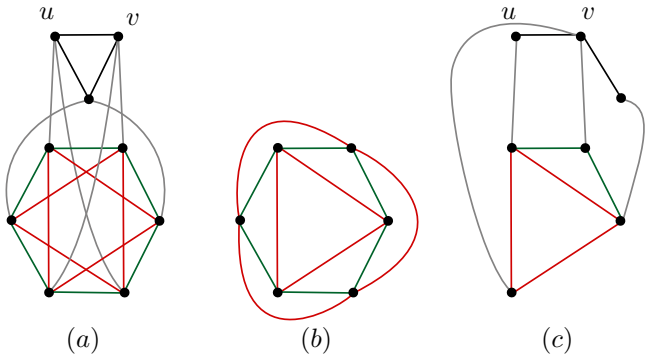

Figure 11. $(9, 21)$

If a red or green edge is deleted, then two of the faces are joined into a 4-gonal face F, which has two opposite vertices x, y in its boundary. Embed the top vertex u that sends edges to x, y inside F, and delete the two remaining top vertices.

Similarly, if a grey edge ux is deleted, we embed u with its one remaining grey edge inside a face of H, and delete the two remaining top vertices.

If a black edge uv is deleted, remove two bottom vertices to obtain a graph embeddable as in Figure 11 (c).

If a grey edge ux is contracted, we are in an easier situation than the previous one: we can delete it along with one more bottom vertex to obtain a sugraph of the previous case.

If a black edge uv is contracted, delete it along with the remaining top vertex to obtain H.

Finally, if a red or green edge xy is contracted, delete it along with a bottom vertex z such that x, y, z are consecutive along the green cycle. We are left with the top triangle, a triangle of H, and a perfect matching joining these triangles. This graph is a triangular prism, hence planar.

6.2.2. *The case* $(11, 22)$. There are five orbits of edges as coloured in Figure 12 (a). Let u, v denote the top vertices.

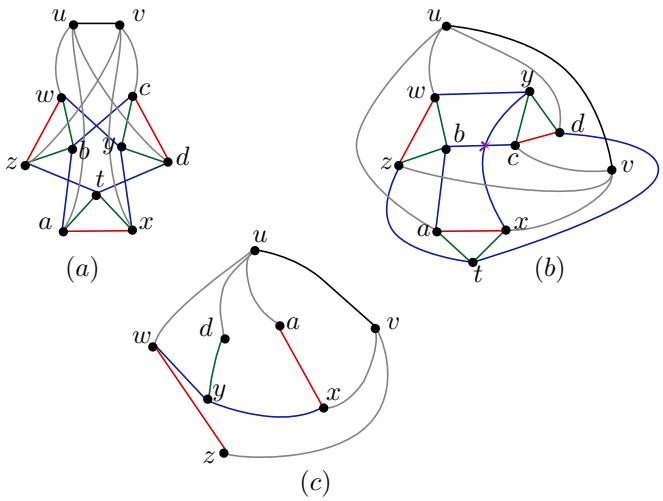

Figure 12. $(11, 22)$

A different drawing, with five crossings of edges, is shown in Figure 12 (b). If the blue xy or grey zv edge is deleted, we remove vertices t, b leaving no crossings. If the black edge uv is deleted, we remove vertices a, x. Notice that if vertices u, v are deleted, then only one crossing remains, namely the one marked with a purple $\times$ symbol. Moreover, if one of the red cd or green yd edges are deleted, then we can reroute the edge xy to avoid any crossing. Thus if a red or green (or blue) edge is deleted, we can delete u, v to obtain a planar graph.

If a red or green edge e is contracted, we remove the resulting vertex as well as the third vertex forming a red-green triangle with e. Since all crossings of Figure 12 (b) involve edges incident with the bottom triangle, this results in a planar graph. If the black uv or grey vz edge is contracted, we remove the resulting vertex as well as the remaining vertex among u, v, z. Then only the crossing marked with a $\times$ symbol remains, and we can re-route the xy edge to avoid it. Finally, if the blue bc edge is contracted, we remove the resulting vertex along with t, and draw the resulting graph as in Figure 12 (c).

6.2.3. *The case* $(10, 21)$. There are four orbits of edges as coloured in Figure 13 (a). The bottom part has a drawing with only one crossing as shown in Figure 13 (b).

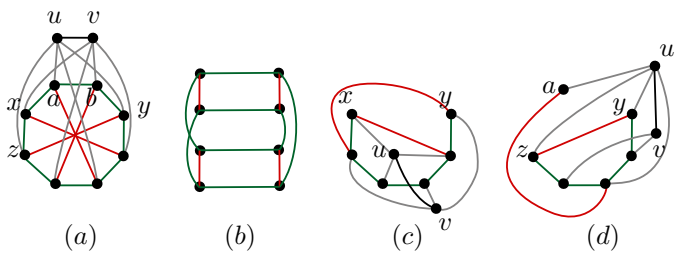

Figure 13. $(10, 21)$

If we delete a green edge, then the aforementioned crossing disappears, and so we obtain a planar subgraph after removing u, v.

Similarly, if a black or grey edge is contracted, then by removing it along with one more vertex we can obtain the subgraph of Figure 13 (b) where one of the vertices involved in the crossing is removed.

Consider now the subgraph of (10,21) obtained by removing a, b, as drawn in Figure 13 (c). There is just one crossing, between the black edge and a green edge. Thus if the green edge ab is contracted, we can delete it along with one of the vertices involved in said crossing to obtain a planar subgraph. If the black edge is deleted, then we remove a, b. If a red edge is contracted, assume it is one incident with a, and remove it along with b; again the crossing of Figure 13 (b) thereby disappears.

For the remaining two cases remove x, b and consider the drawing of the resulting subgraph in Figure 13 (d). If a grey edge is deleted the unique crossing disappears. Finally, if a red edge is deleted, assume it is zy, and embed v at a midpoint of where zy used to lie.

6.2.4. *The case* $(9, 22)_1$. Let u, v denote the top vertices, and note that $G - \{u, v\}$ is a subdivision of K_5.

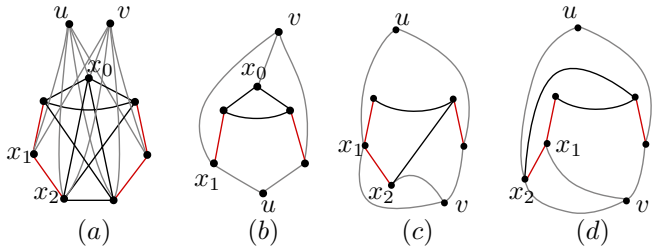

Figure 14. $(9.22)_1$

If one of the four red edges shown in Figure 14 (a) is contracted, then by deleting the contracted vertex as well as its red neighbour we are left with a subgraph of the double suspension of C_4, which is planar.

If an edge ux of u is deleted, we consider two subcases: a) if x is the top vertex x_0, we delete the two bottom vertices. The result is a planar graph shown in Figure 14 (b)

b) if not, then we may assume that x is one of x_1, x_2 because of the symmetry. In this case we delete x_0, as well as z. The result is a planar graph shown in shown in Figure 14 (c) or (d), depending on which of x_1, x_2 is x.

If an edge ux of u is contracted, then we remove v and the contracted vertex. We are left with a proper minor of K_5, hence a planar graph.

By symmetry, if an edge of v is deleted or contracted we are happy.

In all other cases, we remove u, v. We are left with a subdivision of K_5, where the red edges are the ones arising from a subdivision. Thus if we delete any edge, or contract an edge that is not red, we obtain a planar graph.

6.2.5. *The case* $(9, 22)_2$. There are seven orbits of edges as coloured in Figure 15 (a). We embed $G - u$ with four crossings as in Figure 15 (b). If we delete an edge e not incident with u, then we remove u as well as one more vertex to obtain a planar subgraph as follows:

- if $e = ay$ (red), we remove c;
- if $e = xb$ (green), we remove c;
- if $e = xc$ (black), we remove y;
- if $e = cy$ (pink), we remove x;
- if $e = cd$ (blue), we remove x.

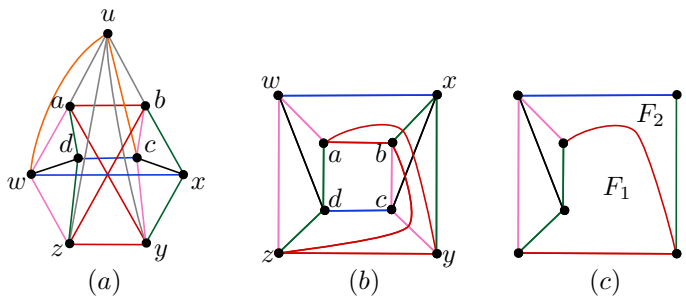

Figure 15. $(9, 21)_2$

If we delete an edge incident with u, say $e = uw$ (respectively, $e = uz$), then we remove vertices c, b, and embed the remaining graph as in Figure 15 (c), putting u inside face F_1 (resp. F_2).

If we contract an edge f, then we remove the contracted vertex and proceed as follows:

- if $f = ab$ (red), we can remove u as Figure 15 (b) becomes planar;
- if $f = xy$ (green), we remove u and obtain a subgraph of the above case $e = xc$;
- if $f = cb$ (pink), or $f = xc$ (black), we remove u and obtain a subgraph of the above case $e = xb$;
- if $f = cd$ (blue), we remove x, and notice that the remaining subgraph is induced by u and a tree, and is therefore planar;
- if $f = uy$ (grey), or $f = uc$ (orange), we remove b;

Acknowledgements. We thank Tâm Nguyễn-Phan for her valuable input and the many discussions we had on embedability of complexes and topology in general. We also gratefully acknowledge the support by the funding by the British Engineering and Physical Sciences Research Council [EP/V009044/1].

References

[1] G. Avramidi and T. Phan. Fungible obstructions to embedding 2-complexes. *arXiv preprint arXiv:2105.10984*, 2021.

[2] R. Daverman. On the scarcity of tame disks in certain mild cells. *Fundamenta Mathematicae*, 1(79):63–77, 1973.

[3] R. J. Daverman. On the absence of tame disks in certain wild cells. In *Geometric Topology: Proceedings of the Geometric Topology Conference held at Park City, Utah, February 19–22, 1974*, pages 142–155. Springer, 2006.

[4] P. de la Harpe. When does a CW-complex of dimension 2 embedd in $\mathbb{R}^4$? MathOverflow. URL: https://mathoverflow.net/q/19618 (version: 2010-03-28).

[5] J. Foisy. A newly recognized intrinsically knotted graph. *Journal of Graph Theory*, 43(3):199–209, 2003.

[6] M. H. Freedman, V. S. Krushkal, and P. Teichner. Van kampen's embedding obstruction is incomplete for 2-complexes in $\mathbb{R}^4$. *Mathematical Research Letters*, 1(2):167–176, 1994.

[7] N. Goldberg, T. W. Mattman, and R. Naimi. Many, many more intrinsically knotted graphs. *Algebraic & Geometric Topology*, 14(3):1801–1823, 2014.

[8] B. Grünbaum. Imbeddings of simplicial complexes. *Comment. Math. Helv*, 44(1):502–513, 1969.

[9] V. Kaluža and M. Tancer. Even maps, the colin de verdiere number and representations of graphs. *Combinatorica*, 42(Suppl 2):1317–1345, 2022.

[10] J. Matoušek, A. Björner, G. M. Ziegler, et al. *Using the Borsuk-Ulam theorem: lectures on topological methods in combinatorics and geometry*, volume 2003. Springer, 2003.

[11] J. Matoušek, M. Tancer, and U. Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *Journal of the European Mathematical Society*, 13(2):259–295, 2010.

[12] M. Pierce. Ytygraphtransforms. https://github.com/mikepierce/YTYGraphTransforms, 2024.

[13] N. Robertson, P. Seymour, and R. Thomas. Sachs' linkless embedding conjecture. *J. Combin. Theory (Series B)*, 64:185–227, 1995.

[14] C. P. Rourke and B. J. Sanderson. *Introduction to piecewise-linear topology*. Springer Science & Business Media, 2012.

[15] H. Sachs. On a spatial analogue of kuratowski's theorem on planar graphs—an open problem. In *Graph Theory: Proceedings of a Conference held in Łagów, Poland, February 10–13, 1981*, pages 230–241. Springer, 2006.

[16] A. Skopenkov. A short exposition of salman parsa's theorems on intrinsic linking and non-realizability. *Discrete & Computational Geometry*, 65:584–585, 2021.

[17] H. van der Holst. Graphs and obstructions in four dimensions. *Journal of Combinatorial Theory, Series B*, 96(3):388–404, 2006.

[18] H. Van der Holst, L. Lovász, A. Schrijver, et al. The colin de verdiere graph parameter. *Graph Theory and Computational Biology (Balatonlelle, 1996)*, 7:29–85, 1999.

[19] H. van der Holst and R. Pendavingh. On a graph property generalizing planarity and flatness. *Comb.*, 29(3):337–361, 2009.

[20] M. Winter. Given a 2-disc embedded in $\mathbb{R}^4$, can I fit another 2-disc with the same boundary? MathOverflow. URL:https://mathoverflow.net/q/406829 (version: 2021-10-22).

Mathematics Institute, University of Warwick, Coventry CV4 7AL, United Kingdom

Email address: a.georgakopoulos@warwick.ac.uk

Email address: martin.h.winter@warwick.ac.uk

Appendix A. Code

We use the Mathematica package `YTYGraphTransforms.m` by Mike Pierce [12] to compute ΔY-families.

A.1. **Enumerating Heawood graphs.** The list of Heawood graphs can be generated as follows:

```
Heawood = WyeTriangleWyeFamily[{
    CompleteGraph[7],
    CompleteGraph[{3, 3, 1, 1}]
}];
Length@Heawood (* output: 78 *)
```

A.2. **Verifying Theorem 6.1.** The following code can be used to verify that each minor of a Heawood graph is 4-flat. The code iterates through all Heawood graphs G and all minors of the form $H \in \{G - e, G/e\}$. It then checks for each pair of vertices $v, w \in V(H)$ whether $H - \{v, w\}$ is planar. If one such pair is found, then H is 4-flat by Theorem 5.9.

```
counterexampleFound = False;
Do[ (* for all Heawood graphs G *)
    minors = Join[
        DeleteDuplicates[
            Table[EdgeDelete[G,e], {e, EdgeList[G]}],
            IsomorphicGraphQ],
        DeleteDuplicates[
            Table[EdgeContract[G,e], {e, EdgeList[G]}],
            IsomorphicGraphQ]
    ];
    Do[ (* for all minors H of G *)
        vertexPairFound = False;
        Do[ (* for each pair of vertices in H *)
            If[PlanarGraphQ[VertexDelete[H, pair]],
                vertexPairFound = True;
                Break[];
            ];,
            {pair, Subsets[VertexList[H], {2}]}
        ]
        If[!vertexPairFound,
            counterexampleFound = True;
            Print["Counterexample found!"];
                (* <-- this code is never reached *)
        ];,
        {H, minors}
    ];,
    {G, Heawood}
]
If[!counterexampleFound,
    Print["No counterexample found!"] (* <-- this code is reached *)
];
(* output: No counterexample found! *)
```

A.3. **The "remaining Heawood graphs".** The seven remaining Heawood graphs (in the sense of Section 6, including K_7 and $K_{3,3,1,1}$) can be listed as follows:

```
remaining = Select[Heawood, Min@VertexDegree[#] > 3 &];
Length@remaining (* output: 7 *)
```