Quartic planar graphs

Jane Tan

October 2018

Australian National University

Declaration

The work in this thesis is my own except where otherwise stated.

Jane Tan

Acknowledgements

First, I'd like to thank my supervisors, Brendan McKay and Scott Morrison. Brendan – your guidance has been invaluable. I'm so grateful for all of the exciting problems you've introduced me to, and that you're still putting up with me a full two years after you agreed to take me on "just for one semester". And Scott – thank you for your help with editing, your calming presence, and for sharing the good chalk.

I would also like to thank all of the lecturers at the MSI who have been so generous with their time and from whom I have learnt so much. Special thanks to Vigleik, Joan and Andrew for teaching fantastic reading courses and supervising me in projects over the years.

Thanks also go to all of the honours students, even the cookie monster and the elusive north-siders. It's been great fun sharing our shiny new office!

Last but not least, thank you to my family for always being on the phone.

Abstract

In this thesis, we explore three problems concerning quartic planar graphs. The first is on recursive structures; we prove generation theorems for several interesting subclasses of quartic planar graphs and their duals, building on previous work which has largely been focussed on the simple or nonplanar cases. The second problem is on the existence of locally self-avoiding Eulerian circuits. As an application of a generation theorem, we prove that all but one 3-connected quartic planar graphs have an Eulerian circuit that is free of subcycles of length 3 or 4. This implies that a 3-connected quartic planar graph admits a P_5-decomposition if and only if it has even order. Finally, we give some new smaller counterexamples to a disproven conjecture of Lovász on circle representations of quartic planar graphs. We also present a gluing construction that tells us a little more about the obstructions to circle representability, although a full characterisation remains elusive.

Contents

Notation

In the following, let G be an undirected graph that may have loops and parallel edges, and v be a vertex of G.

$V(G)$	The set of vertices of G.
$E(G)$	The multiset of edges of G.
$\lvert G \rvert$	The order of G, which is $\lvert V(G) \rvert$.
G^*	The planar dual of a planar graph G.
$d(v)$	The degree of the vertex v.
$N(v)$	The set $\{w \in V(G) : vw \in E(G)\}$ of neighbours of v.
$H \subseteq G$	H is a (not necessarily induced) subgraph of G.
G[A]	The subgraph of G induced by $A \subseteq V(G)$, or $A \subseteq E(G)$.
$G - A$	The subgraph of G obtained by deleting the vertices in a subset $A \subseteq V(G)$ and all incident edges.
$\kappa(G)$	The vertex-connectivity of G.
$\lambda(G)$	The edge-connectivity of G.
P_n	The path on n vertices.
C_n	The cycle on n vertices.
K_n	The complete graph on n vertices.
$K_{m,n}$	The complete bipartite graph with colour classes consisting of m and n vertices.

Chapter 1

Introduction

1.1   Overview

A graph is *k-regular* if every vertex has the same finite degree k. By *cubic* and *quartic*, we mean 3-regular and 4-regular respectively. The former have been well-studied and exhibit many nice properties, so one naturally looks to quartic graphs for interesting extensions of those results, as well as fresh problems that pertain to properties particular to this class.

In this thesis, we explore three problems concerning quartic planar graphs. These have garnered considerable interest as they appear in a wide range of contexts, in large part because there is a nontrivial intersection between the theory of quartic graphs and planar graph theory. For instance, any drawing of a quartic graph gives rise to a quartic planar graph provided no more than two edges cross at any point, by introducing an extra vertex at each crossing. The shadows of knots and links arise in this way (see [37]). Conversely, the face-edge incidences of any planar graph can be encoded as a quartic planar graph.

By focussing on a single class of graphs, we enjoy the opportunity to get acquainted with a particular set of properties, tools and examples that we can take into a variety of problems. This means that although the problems explored may fall under different branches of graph theory, our experiences in one inform our intuition and approach to the others. This has turned out to be quite fruitful. The following highlights from each of the three main chapters are, to the best of our knowledge, all new results.

Our first goal is to provide recursive structures for some interesting subclasses of the quartic planar graphs. These are given in the form of generation theorems, which describe how the graphs in a class can be constructed from some starting

set of graphs by applying a sequence of local expansion operations. We achieve this indirectly by first finding forbidden configuration characterisations of the dual classes of quadrangulations to those that we are primarily interested in generating, and then instead working explicitly with quadrangulations. These results build on the body work on generating simple subclasses of quadrangulations [17] and quartic graphs that are possibly non-planar [27], as well as the class of arbitrary quadrangulations without any restrictions [52].

One of the main motivations for finding recursive structures is that they facilitate inductive proofs. We were able to apply this approach in the context of Eulerian circuits and path decompositions. Quartic graphs are the simplest non-trivial Eulerian graphs, making them a promising starting point for problems involving Eulerian circuits. As our second problem, we consider the restricted class of Eulerian circuits that are free of short subcircuits. We show that all but one 3-connected quartic planar graphs have an Eulerian circuit that avoids 3-cycles and 4-cycles. The proof is constructive, and extends existing work on the avoidance of sub-circuits of length 3 for a similar class of graphs. As a corollary, we also show that a 3-connected quartic planar graph admits a P_5-decomposition if and only if it has even order.

The final problem is in the area of graph drawing, specifically geometric representations of graphs. Here, we consider a disproved conjecture of Lovász that every quartic planar graph can be represented by a system of circles in the plane. We came across this conjecture from a paper by Broersma, Duijvestijn and Göbel [23], who identified it as a potential theoretical application of their graph generation theorem back when it was still open. That method did not work out as hoped, but it did lead us to our work geared toward understanding the obstructions to circle representability, and working toward characterising those graphs that admit a circle representation. On this path, we produce two new families of counterexamples, and improve on the order of the smallest known counterexamples from 822 to 68.

The structure of the thesis is as follows. The first chapter is devoted to recalling some essential notions relating to graph embeddings and connectivity. Each of the next three chapters focusses on an individual problem. We begin by providing some motivation and context for the problem at hand, before surveying recent progress in the area and proving our main results. A list of open problems and ideas for future directions are also included for each one.

1.2 Conventions

We will assume basic familiarity with graph theory. For a refresher, we refer to Diestel [26] with a warning that we occasionally deviate from the notation and definitions used there. In this thesis, we use the following conventions.

A *graph* is a pair $G = (V, E)$ such that V is a finite set, and E is a finite multiset consisting of elements of $[V]^2 \cup \{(v, v) \mid v \in V\}$ where $[V]^2$ denotes the set of all 2-element subsets of V. An edge that has both endvertices being the same vertex is called a *loop*, and two edges are said to be *parallel* if they have the same endvertices. By this definition, our graphs are assumed to be finite, undirected, and unless otherwise indicated may have loops and parallel edges. In addition, we will assume that our graphs are connected unless otherwise indicated.

We say that an edge is *incident* to its endvertices, and vice versa. Two vertices are *adjacent* if there is an edge between them, meaning incident to them both. The degree of a vertex v, denoted by $d(v)$ is the number of edges incident to v, with loops counted twice. If an edge e is not parallel to any other edges, we say that e is *simple*. A *simple graph* is then a graph that is loopless and in which every edge is simple. We sometimes use the term *multigraph* to emphasise that loops and the parallel edges may be present, but this should be read no differently than 'graph' as we have defined it above.* We refer to any subgraph consisting of only two parallel edges and their endvertices as a *digon*.

A *walk* of length k between v_0 and v_k in a graph G is an alternating sequence of vertices and edges $v_0 e_0 v_1 e_1 \ldots e_{k-1} v_k$, where each two consecutive objects are incident. Here, v_0 and v_k are endvertices of the path, whilst v_i for $1 \leq i \leq k - 1$ are the inner vertices. Walks are not directed, but it is convenient to say that they 'start' and 'end' at particular vertices so we will frequently do so.

There are a few different notations to denote walks that we shall use depending on context. For simple graphs, since each edge is uniquely identified by its endvertices it is enough to write down the sequence of vertices in the form $v_0 v_1 \ldots v_k$, so that the implicit sequence of edges is $v_0 v_1, v_1 v_2, \ldots, v_{k-1} v_k$. Long walks, such as those in Chapter 4, are sometimes abbreviated to a form such as $v_0 X v_k$ in which X represents the walk $v_1 \ldots v_{k-1}$ where v_0 is adjacent to v_1 and v_{k-1} is adjacent to v_k. When parallel edges exist, we must be careful to distinguish which edge we are taking between two vertices. In that case we include the edge names as well. For sake of readability we adopt the format $v_0 \overset{e_0}{—} v_1 \overset{e_1}{—} v_2 \ldots v_{k-1} \overset{e_{k-1}}{—} v_k$, which

---

*The term *multigraph* (also, *pseudograph*) is sometimes used to denote what we are calling a graph, whilst a *graph* elsewhere is often assumed to be simple.

we will see regularly in Chapter 3.

We say that a walk is *closed* if its final vertex is the same as its initial vertex, and *open* otherwise. A *circuit* is then a closed walk of length at least 2 in which each edge occurs at most once. An *Eulerian circuit* is a circuit in which every edge is traversed precisely once. A walk in which each vertex occurs at most once (hence each edge also appears at most once), except possibly where the initial vertex is the same as the last, is called a *path*, and a closed path of length at least two is called a *cycle*. A cycle with an even number of edges is an *even* cycle, and *odd* otherwise. We use the term *digon* to refer to a cycle of length 2.

Example 1.1. Two graphs are shown in Figure 1.1. The first graph, G, is not simple because e_6 is a loop, e_3 is parallel to e_3', and e_5, e_5' and e_5'' are all parallel. The walk $d\overset{e_3}{\text{---}}c\overset{e_3'}{\text{---}}d\overset{e_4}{\text{---}}a\overset{e_6}{\text{---}}a\overset{e_1}{\text{---}}b\overset{e_2}{\text{---}}c\overset{e_3}{\text{---}}d$ is closed but not a circuit since e_3 is repeated, and it contains a circuit $d\overset{e_4}{\text{---}}a\overset{e_6}{\text{---}}a\overset{e_1}{\text{---}}b\overset{e_2}{\text{---}}c\overset{e_3}{\text{---}}d$ that is not a cycle since a is repeated. The walk $be_5de_5'b$ is a digon, as is $be_5''de_5'b$.

On the other hand, H is a simple graph. The cycle $ae_1be_2ce_3de_4a$ can be abbreviated to $abcda$. An example of a closed walk that is not a circuit is $abdcba$.
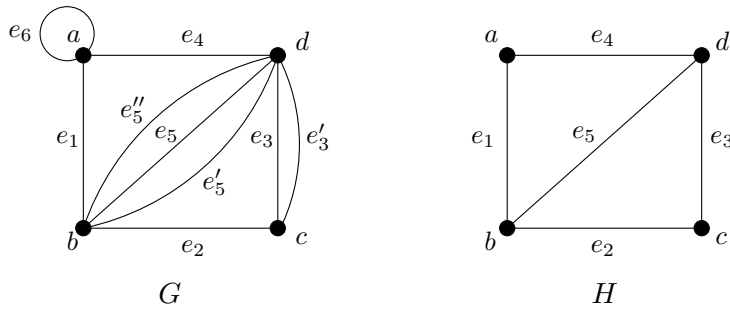

Figure 1.1: A graph and a simple graph.

There are certain special graphs and families of graphs with widely accepted names in circulation. In case an unfamiliar name comes up, we invite the reader to visit the gallery in Appendix A in which we display the names, descriptions and some pictures of all such graphs that we refer to in this thesis.

Chapter 2

Toolkit

In this chapter, we record some standard results and constructions in graph theory, as well as properties specific to quartic planar graphs and the closely related class of quadrangulations. These results are mostly classical, to the extent that their proofs in literature tend to be omitted or else invoke folklore[*]. In those cases, we supply a short proof or reference.

2.1 Plane and planar graphs

Definition 2.1. An embedding of a graph G in the plane is an injective map that sends $V(G)$ to points in the plane and edges of $E(G)$ to disjoint simple open arcs such that the image of each edge joins the images of its endvertices.

Such an embedding corresponds to a drawing of the graph in the plane so that there are no edge crossings. The leads to the following terms that we should be careful to distinguish.

Definition 2.2. A graph is *planar* if it has an embedding into the plane. The image is called a *plane* graph, which contains the data of an abstract graph together with a specified embedding.

The faces of a plane graph G are the maximal connected subsets of the open set $\mathbb{R}^2 - G$. Note that each of these is homeomorphic to an open disk, so we say that this is a *cellular* embedding. This is automatic for connected graphs drawn on the plane, but not necessarily so if we embed graphs on other surfaces. The topological boundary of any face corresponds to a walk in the graph, and we say

---

[*]This is quoted from the proof of Lemma 1 in [17].

that the face is incident to each of the edges and vertices in its boundary. More generally, by the Jordan Curve Theorem, any cycle in a graph bounds exactly two regions. There is always one unbounded face which we call the outer face, and similarly one unbounded region which we call the outer region or exterior, as opposed to inner regions and interior of a cycle. If at least one of the regions bounded by a particular cycle is also a face, then we say that the cycle is *facial*, otherwise it is *non-facial* or *separating*.

Remark 2.3. Having an embedding in the plane is equivalent to having an embedding in the sphere S^2, since this is the one point compactification of the plane. Given an embedding of a graph on the sphere, we may assume without loss of generality that no part of the graph lies on the north pole by possibly rotating. The stereographic projection then gives a correspondence between spherical and planar embeddings.

One should think of the difference as being that a planar embedding comes with a particular choice of outer face, which corresponds on the sphere to the face that contains the north pole. We will freely switch between these two perspectives, as although our diagrams must necessarily be planar, it is advantageous to think of our graphs as being on the sphere since this means all faces are treated equally.

In practice, when we fix an embedding on the plane as opposed to the sphere, the only difference is that the planar embedding comes with a choice of outer face so it then makes sense to refer to the interior and exterior regions of any cycle. This is useful, for instance, in the following situation.

Lemma 2.4. *A plane graph G for which every face is bounded by a walk of even length is bipartite.*

Proof. Fix a planar embedding of G. We may assume that every edge in G lies on the boundary walks of two distinct faces by possibly adding parallel edges. Any cycle in G can be viewed as a finite union of the faces in its interior. Suppose we count the edges inside the cycle by counting the edges adjacent to each face; any edge not lying in the cycle will be counted twice, one from each face adjacent to it, whilst edges on the cycle are counted once. Thus, the length of the cycle is equal to the sum of the lengths of each face modulo 2. Since each face is even by assumption, it follows that G cannot have any odd cycles. Therefore, G is bipartite. □

The plane graphs in which every face is bounded by a walk of the same even length form a special class of graphs that satisfy the conditions of the previous lemma. More generally:

Definition 2.5. A k-angulation (of the sphere) is a plane graph in which every face is bounded by a walk of length k for $k \geq 3$.

The 3-, 4- and 5-angulations are called triangulations, quadrangulations and pentangulations respectively.

Corollary 2.6. *Quadrangulations are bipartite.*

Provisionally, we say that two plane graphs are equivalent if they are related by a homeomorphism of the sphere that restricts to an isomorphism of graphs, and regard an embedding as an equivalence class. Observe that a graph embedded on the plane can be transformed into an equivalent one with a different choice of outer face by applying a Möbius transformation. A second equivalent definition will be given once we introduce combinatorial embeddings. It is certainly possible for a graph to have inequivalent embeddings. An example is shown in Figure 2.1.
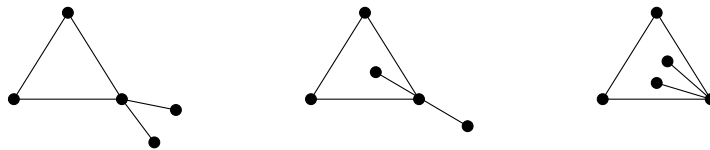

Figure 2.1: Three planar embeddings of a planar graph; the first is equivalent to the third but not to the second.

There are two famous characterisations of simple planar graphs. Recall that *subdividing* an edge in a graph means to replace the edge by a path between its endvertices such that each inner vertex has degree 2. An edge *contraction* is performed by removing an edge xy together with all incident edges, and identifying x and y into a new vertex that is incident to all of the edges formerly incident to x or y. It is implicit that any new parallel edges created are also merged. A graph G' is a *subdivision* of G if it can be obtained by subdividing edges of G, and it is a *minor* of G if it can be constructed from G by a series of edge contractions, vertex deletions and edge deletions. An example is shown in Figure 2.2.

Theorem 2.7 (Kuratowski). *A simple graph is planar if and only if does not contain a subdivision isomorphic to K_5 or $K_{3,3}$.*

Theorem 2.8 (Wagner). *A simple graph is planar if and only if it contains neither K_5 or $K_{3,3}$ as a minor.*
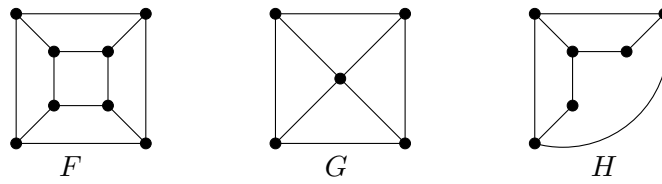
Figure 2.2: The graph F contains G as a minor and a subdivision of H.

The following result of Euler is classical. For now, we only state the simplest version for the plane.

Theorem 2.9 (Euler's formula). *For any plane graph with n vertices, e edges and f faces,*

$$n - e + f = 2.$$

Since the number of faces is determined by the number of edges and vertices, we can refer to the number of faces of a planar (rather than plane) graph as a well-defined quantity. Typically, Euler's formula is applied to show that certain graphs are non-planar, or to obtain bounds on the number of vertices, edges or faces given some relationships between these values. For example, the following is standard.

Proposition 2.10. *There are no simple planar k-regular graphs for $k \geq 6$.*

Proof. Let G be any simple planar graph with n vertices, e edges and f faces. We may assume that $v \geq 3$ and $e \geq 3$, so each face is bounded by at least three edges. Then $3f \leq 2e$. Substituting this into Euler's formula gives $2e - 3e \geq 6 - 3n$ and hence $2e \leq 6n - 12$. Diving through by n, we find that the average degree is at most $6 - \frac{12}{n}$ which is strictly less than 6. Therefore, G must have a vertex of degree at most 5. □

Although we are primarily concerned with graphs embedded on the sphere, one can also embed graphs onto other orientable and non-orientable surfaces of higher genus. Our surfaces are specifically compact 2-manifolds without boundary, and we consider only cellular embeddings. An example of an embedding that is not cellular is shown in Figure 2.3. The term *map* is used as a generalisation of a plane graph, and refers to an abstract graph together with an embedding into some surface. Just as not all graphs are planar, it is also the case that only certain graphs can be embedded in certain surfaces. The famous graph minor theorem due to Robertson and Seymour [73] implies that there is an analogue of Wagner's theorem for each surface. We also have a more general version of Euler's formula.
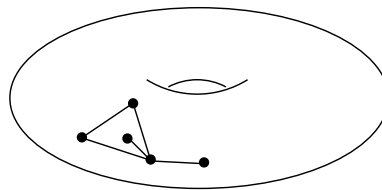
Figure 2.3: A non-cellular embedding of a graph on the torus.

Theorem 2.11 (Euler's formula)**.** *For any graph with n vertices, e edges and f faces embedded on a surface* Σ,

$$n - e + f = \begin{cases} 2 - 2g, & \textit{if } \Sigma \textit{ is orientable with genus } g, \\ 2 - h, & \textit{if } \Sigma \textit{ is non-orientable with genus } h. \end{cases}$$

2.1.1 Combinatorial embeddings

The definition we have so far of an embedding is specifically a topological embedding. There is also a purely combinatorial definition, which is often more convenient to work with.

Definition 2.12. A *rotation* at a vertex in a map is the cyclic ordering of edges incident to that vertex induced by a specified orientation. A *rotation system* for a graph consists of a choice of rotation at each vertex.

There is a one-to-one correspondence between cellular embeddings on orientable surfaces and the possible rotation systems of a graph, a fact which is sometimes called the Heffter-Edmonds-Ringel rotation principle [61]. That is, any cellular embedding of a plane graph gives rise to a unique rotation system, and conversely specifying the cyclic order at each vertex determines a cellular embedding of the graph into some orientable surface. For the proof of this correspondence, we refer the reader to [39, 61].

Note that given a graph and a rotation system, it is straightforward to deduce the number of faces of the cellularly embedded graph to which it corresponds. To do this, we trace out the boundary cycles of a face by choosing an initial vertex together with an initial incident edge and choice of "left" or "right", and then turning sharp left or sharp right according to the chosen direction at each vertex until we return to the initial vertex and initial edge in the correct orientation. Repeating this process until all edges have been traversed in both directions gives the collection of faces. One can then determine the genus of the embedding surface using Theorem 2.11.

There is also a twisted version of a rotation system that can be used to encode graphs embedded on non-orientable surfaces which is detailed in [39, 61].

The following notion of equivalence for combinatorial embeddings corresponds to our earlier topological definition.

Definition 2.13. Two maps are regarded as equivalent if they are related by an isomorphism of abstract graphs that also either preserves the cyclic ordering of edges at each vertex, or reverses the ordering at every vertex.

2.1.2   Dual, medial and radial graph constructions

We now record several related constructions on plane graphs that encode information about the incidences between vertices, edges and faces. Recall that by default, we assume that all graphs are connected.

Planar duals

To each planar graph G, we can associate another planar graph G^* called its *dual* as follows. Place a vertex in each face of G, so $V(G^*)$ is the set of faces of G. For each edge e in G, create an edge e^* in $E(G^*)$ that crosses e with endvertices corresponding to the two faces adjacent to e. If e is only adjacent to only one face, then we get a loop in the dual graph. Evidently, the dual of a quartic plane graph is a quadrangulation.

Since abstract planar graphs may have inequivalent embeddings, a planar graph can have nonisomorphic duals. However, it is true that each plane graph has a unique dual. This means that a class $\mathcal{C}$ of plane graphs is in bijective correspondence to the dual class $\mathcal{C}^* := \{G^* \mid G \in \mathcal{C}\}$.

The construction we have described is of the *geometric dual* graph. There is a separate notion of *combinatorial dual* which has quite a different description but is equivalent.

Medial graphs

The medial graph MG of a plane graph G is a quartic plane graph that encodes the edge-face incidences of G. To construct this, we let $V(MG)$ correspond to edges of G, and create an edge between two vertices for each face of G on which these vertices occur consecutively in the boundary. A plane graph and its dual always have the same medial graph.

In fact, every 4-regular plane graph is the medial graph MG of some plane graph G. We call the associated graph a Tait graph, as this construction was originally used by Tait in the study of knots and links. The proof is by construction. Given a quartic plane graph G, create a checkerboard colouring of its faces with two colours, say black and white, so that each edge is incident to one black face and one white face. Let us specify that the outer face is white. This is possible since G is Eulerian, so G^* is bipartite by Lemma 2.4 and a proper 2-vertex-colouring of G^* translates to the desired checkerboard colouring of G. Now let $V(TG)$ correspond to the black faces, and the edge set consist of xy whenever the faces corresponding to x and y share a vertex. Choosing the white faces gives $(TG)^*$, which also satisfies $MTG \sim G$. These are the only two embedded graphs up to isomorphism that have G as their medial graph.

Radial graphs

Given a plane graph G, we define the radial graph RG to be the quadrangulation $(MG)^*$. This can also be constructed directly. Let $V(RG)$ be the vertices and faces of G, with an edge $xy \in E(RG)$ whenever x and y correspond to a vertex and one of its incident faces, counting multiplicities. Evidently, the radial graph encodes vertex-face incidences of G.

Conversely, from each quadrangulation we can construct a graph DG by properly colouring the vertices of G with two colours, which is possible by Lemma 2.6, and taking the vertices of one colour class as $V(DG)$. Two vertices are adjacent whenever they lie on the same face of G. Again, using the other colour class gives $(DG)^*$ which has the same radial graph G as DG. There does not appear to be a standard name for this construction, so we use D for "diagonal" since the new edges are diagonals of the quadrangular faces of G.

These constructions are illustrated in Figure 2.4. In going from the medial and radial graphs to the central plane graphs, taking the black colour class yields the lower middle graph, whilst its dual is obtained by taking the white colour class. As a shorthand, we will sometimes speak of the Tait graph TG of a planar graph G (as opposed to a plane graph), in which we implicitly fix some plane embedding and, where required, a choice of colouring to make it well-defined. We only do this when the choices do not matter.
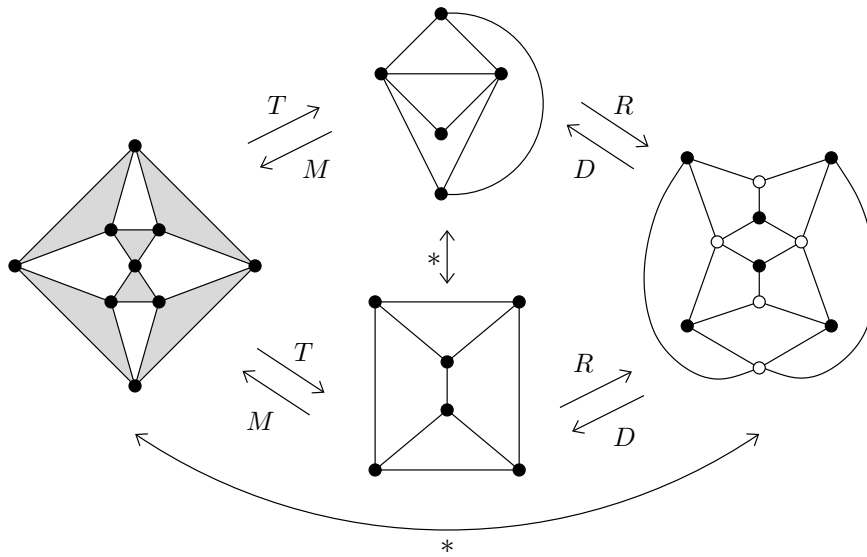
Figure 2.4: Dual plane graphs in the centre, together with their medial graph on the left and their radial graph on the right.

2.2   A note on interpreting diagrams

Although planar graphs, by definition, are a nice class of graphs to draw, in the wild we almost never have the luxury of seeing the entire graph. Rather, it is common to work locally in small neighbourhoods of a graph at a time, but also have some knowledge about where the rest of the graph lies in terms of rotations at particular vertices. This poses some challenges when we wish to display such information precisely using diagrams.

The pictures in this thesis all represent plane (embedded) graphs, and so the cyclic ordering of edges at each vertex is important. It should be noted that the embedding shown may not be the unique possible embedding, but this will be made clear whenever it is important. To represent general structures diagrammatically, we adopt a system of edge placeholders consisting of half-edges, black triangles and white triangles which may be interpreted according to the following rules.

(i) All vertices shown are distinct from each other.

(ii) Edges and half-edges (which indicate an edge must be present but one end-vertex is not depicted) occur in the cyclic order in which they are shown at each vertex.

(iii) A white triangle at a vertex indicates that zero or more edges are present in this position.

(iv) A white triangle with an extra half-edge at a vertex indicates that at least one edge is present in this position. A bounded region containing this placeholder cannot be a facial.

(v) For any given region that contains black triangles at each incident vertex, that region must not be a face of the graph; in other words, at least one of those black triangles associated with the particular region is actually realised as one or more edges.

(vi) Any two edges that are shown consecutively in the cyclic order (not separated by any edge placeholder) at a vertex must follow each other directly in the cyclic ordering at that vertex. If a region contains no edge placeholders, then it is necessarily facial.

By convention, facial cycles will be written clockwise except the outside face which will be written anticlockwise.
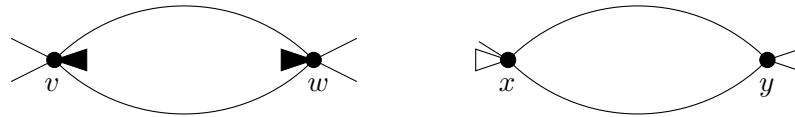

Figure 2.5: Applying the above conventions, the left digon is non-facial due to the black triangles on the inside and half-edges on the outside, whereas the right digon is facial since there is no placeholder inside the parallel edges. The region bounded outside is not a face because one of the white triangles has a half-edge.

2.3 Connectivity

Recall that a non-empty graph is *connected* if there is a path between any two of its vertices, and *disconnected* otherwise, in which case we refer to the maximal connected subgraphs as *components*.

Definition 2.14. A k-vertex-cutset (or k-vertex-cut) in a graph G is subset $S \subseteq V(G)$ with $|S| = k$ such that $G - S$ is disconnected. The unique element in any 1-vertex-cut is called a *cutvertex* of the graph. A k-edge-cutset (or k-edge-cut) in a graph G is a subset $T \subseteq E(G)$ with $|T| = k$ such that $G - T$ is disconnected. The unique element in any 1-edge-cut is called a *bridge*.

Definition 2.15. A graph is k-vertex-connected if it does not have any cutsets with fewer than k vertices, and k-edge-connected if it does not have any edge-cutsets with fewer than k edges.

The greatest integer k such that a graph G is k-vertex-connected is the *vertex connectivity* of G, denoted by $\kappa(G)$. Similarly, the greatest integer k such that a graph G is k-edge-connected is the *edge connectivity* of G, denoted by $\lambda(G)$. When there is little chance of confusion, we sometimes drop the 'vertex' part of the above terms and just write, for example, k-connected to mean k-vertex-connected.

This definition measures how connected a graph is by how hard it is to disconnect. In some ways, Menger's theorem provides a more intuitive description.

Theorem 2.16 (Menger)**.** *A graph is k-edge-connected if and only if there are k edge-disjoint paths between any two vertices. Similarly, a graph is k-vertex-connected if and only if there are k internally vertex-disjoint paths between any two vertices.*

Planar graphs with sufficiently high connectivity have several nice features.

Proposition 2.17. *A simple plane graph other than K_1 and K_2 is 2-vertex-connected if and only if every face is bounded by a cycle.*

Proof. See 4.2.6 in [26]. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $\square$

Lemma 2.18. *Every 3-vertex-connected quadrangulation is simple.*

Proof. Quadrangulations are bipartite by Lemma 2.6, and hence loopless. By definition they cannot have facial digons, so if any parallel edges exist, their endvertices would form a 2-cut. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $\square$

Theorem 2.19 (Whitney)**.** *The 3-connected planar graphs have unique embeddings on the sphere.*

The last theorem provides a reason for why 3-vertex-connected planar graphs are of special interest, and imposing this condition gives a further refinement of the quartic planar graphs that makes for a good entry point when approaching difficult problems.

There is one other type of connectivity that we will meet.

Definition 2.20. A simple planar graph is *cyclically k-edge-connected* if no two circuits of G can be separated by removing fewer than k edges. If, in addition, whenever two circuits of G are separated by the removal of k edges then one of the components is a simple cycle, we say that G is *exactly cyclically n-edge-connected.*

Equivalently, a graph G is exactly cyclically k-edge-connected if it is cyclically k-edge-connected but not cyclically $(k+1)$-edge-connected. Interest in cyclic connectivity arose in connection to the 4-colour theorem, with one famous reduction due to Birkhoff being that it is sufficient to prove the 4-colour theorem for exactly cyclically 5-connected planar graphs [69].

These notions of connectivity are interdependent. Being 1-vertex-connected, 1-edge-connected or connected are equivalent. It is straightforward to see that any graph with a k-edge-cutset also has a k-vertex cutset by taking one endvertex of each of the edges in the cut. This means that if G is k-vertex-connected then it is also k-edge-connected. That is, for any non-trivial graph G we have $\kappa(G) \leq \lambda(G) \leq \delta(G)$ where $\delta(G)$ is used here to denote the minimum degree of any vertex in G. Similarly, a graph that is cyclically k-edge-connected is k-edge-connected. The following is another easy deduction.

Proposition 2.21. *If G is cubic, then $\kappa(G) = \lambda(G)$.*

Further relationships exist for planar graphs. One that will be useful for us when we later consider generating triangulations is:

Proposition 2.22. *A triangulation is n-vertex-connected if and only if its dual is cyclically n-edge-connected.*

2.3.1 Cuts in quartic planar graphs

From the construction described earlier it is clear the dual of a simple graph need not be simple, and the reason for this is secretly encoded in edge-cutsets. The following classical result makes this precise.

Lemma 2.23 (Cut-cycle duality). *If G is a connected plane graph, there is a one-to-one correspondence between cycles of G and minimal edge-cutsets of G^*.*

This is Theorem 4.6.1 in [26]. Importantly, this tells us that the girth of a graph, which is the length of the shortest cycle, is the edge-connectivity of the dual graph. Hence, if a planar graph G is connected but not bridgeless, then G^* is a multigraph with loops and possible parallel edges. If G is bridgeless but not 3-edge-connected, then G^* is loopless but has parallel edges. Finally, if G is 3-edge-connected, then G^* is simple. We will go further with determining similar dual structures in Section 3.2.2.

The following is an immediate but noteworthy consequence.

Proposition 2.24. *There are no cuts consisting of an odd numbers of edges in a quartic planar graph.*

Proof. We saw in Lemma 2.6 that quadrangulations of the sphere are bipartite, and hence do not have any odd cycles. Therefore their duals, which are quartic plane graphs, do not have odd cutsets by Lemma 2.23. It follows that this is also true of quartic planar graphs by choosing any planar embedding. □

For quartic plane graphs, we can also say very explicitly what 1-vertex-cuts and 2-vertex-cuts look like. This is useful in cases where we would like to decompose a graph into blocks that have higher connectivity than the whole graph in order to apply some property that normally required 3-vertex-connectedness. The characterisation tells us how blocks might be joined together.

Let G be a quartic planar graph with a cutvertex v. This cut can only take one of two forms, as shown in Figure 2.6, corresponding to the situations where v is adjacent to two edges in each component of $G^* - v$, or three in one component and one in the other. If the configuration on the right of Figure 2.6 is present, then since G^* is quartic, each component of $G^* - v$ has an odd number of vertices with odd degree which is impossible. Thus, any 1-cut takes the left configuration.
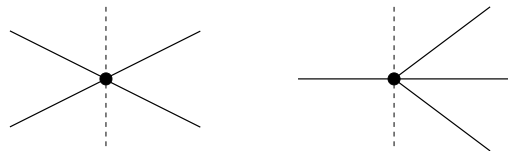

Figure 2.6: Possible 1-cut configurations

A similar analysis shows that there are four 2-vertex-cut configurations. Suppose G has a 2-cut $\{x, y\}$. It is possible that $G - \{x, y\}$ has more than two components, but we may assume that the cut has two *sides* which need not be connected. An identical argument to the 1-cut case shows that there must be an even number of edges between $\{x, y\}$ and each side. If $xy \in E(G)$, this together with the fact that G is bridgeless (Proposition 2.24) implies that x and y each have one neighbour in, say, the left side of the cut, and two each in the right side. If x and y are not neighbours, then x either has two neighbours on each side and hence y must as well, or else x has one neighbour on the left side and three on

the right, in which case y has one on the left and three on the right, or 3 on the left and one on the right. These are all the possibilities up to symmetry.

We will typically represent these as shown in Figure 2.7, where the shaded regions represents the sides of the cut. The endvertices of the half-edges shown may not be distinct. The top row are *balanced* cuts whilst the other three are *unbalanced.* Observe that the unbalanced 2-cuts are all related; the first one can be seen as a special case of the second, and any graph that has the second type of 2-cut must also have the third and vice versa by replacing one of the cut vertices with its unique neighbour on one side of the cut.
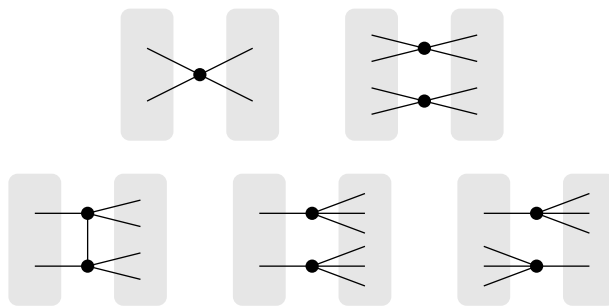

Figure 2.7: All 1- and 2-cut configurations

Chapter 3

Recursive generation

Informally, the idea of generating a class of graphs is that starting with a subclass of basic graphs which are in some way simpler, we can recursively apply local expansion operations to gradually increase their complexity, as measured by some parameter. The superclass of graphs that can be obtained in this manner is said to be generated from the starting set, and the corresponding result is a *generation theorem* for the larger class.

Apart from being insightful structural results, the appeal of having such descriptions for particular classes of graphs stems from two main applications. First, generation theorems facilitate inductive proofs. If one can show that some property is satisfied by the starting graphs and preserved by the expansion operations, then it would follow that it is also satisfied by graphs in the superclass. The second application is to practical graph generators, where implementations of the theorem make it possible to enumerate the graphs in the class (up to some parameter) and – in contrast to theoretical enumeration – output an exhaustive list. This is useful for checking conjectures, and also has applications to chemistry where certain molecular structures can be generated (see for instance [15, 16, 35, 40]).

The main result in this chapter is a collection of generation theorems of quadrangulations and their dual classes of quartic graphs, together with the proof of those dual relationships. In light of these characterisations of the dual classes, we will instead explicitly be generating quadrangulations as we find this to be a more convenient setting in which to work. We will begin with a more abstract discussion of recursive generation, and familiarise ourselves with generation theorems for some closely related classes of graphs.

3.1   Definitions and literature

The following key definition formalises our earlier description of what it means to recursively generate a class of graphs.

Definition 3.1. Given a class of graphs Y and a subclass X together with a set of operations, we say that the class Y can be generated from class X by those operations if for every graph in $G \in Y$ there exists a sequence of graphs $G_0, G_1, \ldots, G_n = G$ such that $G_0 \in X$, $G_i \in Y$ for each $1 \leq i \leq n$ and G_{i+1} can be obtained from G_i by applying one of the operations.

That is, one can obtain any graph in Y by recursively applying operations starting with some graph in X. We are specifically interested in those operations that increase some parameter related to the graph. Some common possibilities here include the number of vertices, edges or faces, as well as possibly multigraph structures. These operations are typically called *expansions*, so accordingly their inverses are *reductions*. Note that the class we are generating need not be closed under the set of expansion operations, but it is important that after applying a reduction to any graph in the class, the resulting graph is still in the class.

With a view to practical applications, there are some general qualities that we like to have in a generation theorem. It is usually the case that the starting class of graphs to either be a small finite collection graphs, or else some infinite family for which there is a general, explicit description such as the cycles, prisms, or pseudo-double wheels. It is also common for generation theorems to build upon others. For instance, if we know how to generate some class B of graphs from A, then it may be convenient to start with B to generate some larger superclass C. Of course, this could then be phrased as a statement generating C from A by simply concatenating the lists of operations. Regarding choice of expansion operations, it is generally desirable to have as few operations as possible, and with each one being as restricted as possible, meaning those for which the inverse reduction operations can be applied in fewer places are considered more efficient.

Other methods of practical generation

The *reduction-expansion* generation theorems that we are discussing are not the only kind being studied. There are, broadly, two other main approaches. The first is still recursive and uses local operations, but they are no longer required to increase some parameter. A simple example of such an operation for triangulations is a diagonal flip, which involves taking a 4-cycle *abcd* drawn so that it

has only one edge, say *ac*, inside, and replacing that edge with the other diagonal *bd*. A classical theorem of Wagner [86] is that any two triangulations are related by a sequence of diagonal flips. Analogous diagonal transformations and results exist for quadrangulations [64] and pentangulations [51].

The second method uses patches, which are subgraphs obtained by cutting a graph into parts via some well-defined paths from which it can be assembled. This has been used in the generation of fullerenes [35], as well as to generate cubic and quartic planar maps with prescribed numbers of vertices and face degrees [18].

The advantage of reduction-expansion generation is that it is best suited for the applications we have mentioned, which is clearly evident for induction arguments in particular since we would hope for an easy base case. Henceforth, by generation theorem we always mean the reduction-expansion sort.

3.1.1 Generating planar graphs

Polyhedra, which are the 1-skeleta of convex 3-polytopes, were among the first classes of graphs for which recursive structures were proved (see [30, 40, 77]). A classical theorem of Steinitz from the 1930s says that the 3-connected planar graphs are all, up to isomorphism, the graphs of convex polyhedra. The original proof is by induction, and one version of the statement is a generation theorem for planar 3-connected graphs.

Theorem 3.2 (Steinitz, [77])**.** *The planar 3-connected graphs can be generated from the tetrahedron (K_4) by operations known as face splittings.*

There is a wealth of literature on generation theorems, and on generating maps in particular. For the most part, we will focus on *k*-regular graphs and *k*-angulations of the sphere. We have chosen to phrase each of the generation theorems in the world of *k*-angulations, since this is popular in literature and convenient for our proofs.

Triangulations

The natural starting point for problems concerning planar graphs is to consider the triangulations (of the sphere), since any planar graph can be obtained by deleting edges from a triangulation of the same order.

The first result we record is the dual of Theorem 3.2. We did not define a face splitting before, but its dual, known as vertex splitting, is shown in Figure 3.1.
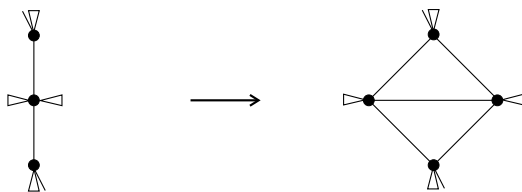
Figure 3.1: Vertex splitting.

Theorem 3.3 (Steinitz, [77]). *The class of simple triangulations can be generated from the tetrahedron (K_4) by vertex splittings.*

The next set of expansions, denoted by O_k for $k = 3, 4, 5$, are attributed to Eberhard [30], and defined as follows: for a k-cycle $x_1 x_2 \ldots x_k x_1$ drawn so that there are $k - 2$ edges in its interior, delete those inner edges and create a new vertex u inside the cycle together with the edges ux_i for $i = 1, 2, \ldots k$. These give an alternative generation theorem for the class of triangulations, which is more efficient than the last theorem because their inverses require more specific local configurations to be applied.

Theorem 3.4 (Bowen and Fisk [13]). *The class of simple triangulations can be generated from the tetrahedron by O_3, O_4 and O_5.*

Using this result, Bowen and Fisk were able to enumerate all of the nonisomorphic simple triangulations with up to 12 vertices [13]. Note that simple triangulations are all 3-connected, but O_3 creates separating 3-cycles and in particular 3-vertex-cutsets, and so clearly cannot be used to generated 4-connected triangulations. Perhaps surprisingly, the other two operations are sufficient.

Theorem 3.5 (Brinkmann, Larson, Souffriau, and Van Cleemput [19]). *The class of all 4-connected triangulations can be generated from the octahedron by the operations O_4 and O_5.*

Equivalently, this generates the dual class of cyclically 4-connected cubic plane graphs. Other restricted classes for which generation theorems have been published are the 3-connected triangulations with minimum degree 5 (see [5], corrected in [21]), 4-connected triangulations with minimum degree 5 [21], 5-connected triangulations [4, 24], 5-connected triangulations with minimum degree 4 [7], and Eulerian triangulations [7]. In the dual form, these classes can be tidily expressed using cyclic connectivity and Lemma 2.22 (see also [21]).

Quadrangulations

There are a basic set of expansion operations shown in Figure 3.2, small variations of which account for most generation theorems of classes of quadrangulations.

To get an overview, we start with the following main theorems in [17], in which the authors improved a number of existing theorems for generating simple quadrangulations.

Theorem 3.6 (Brinkmann, Greenberg, Greenhill, McKay, Thomas, Wollan [17])**.** *Let the expansion operations be as labelled in Figure 3.2.*

1. *The class of all simple quadrangulations of the sphere can be generated from the square (C_4) by the expansions E_0 and E_1.*

2. *The class of all simple quadrangulations of the sphere with minimum degree 3 can be generated from the pseudo-double wheels by E_1 and E_3.*

3. *The class of all 3-connected quadrangulations of the sphere can be generated from the pseudo-double wheels by E_1 and E_3.*

4. *The class of all 3-connected quadrangulations of the sphere in which all 4-cycles are facial is generated from the pseudo-double wheels by E_1.*

One reason that these classes of graphs are interesting is that they are precisely the classes of radial graphs of the loopless 2-connected (but not necessarily simple) plane graphs, loopless 2-connected plane graphs without no facial digons nor vertices of degree 2, simple 2-vertex-connected and 3-edge-connected plane graphs, and simple 3-connected plane graphs respectively. In addition, their dual classes are the 4-edge-connected quartic plane graphs (not necessarily simple), simple 4-edge-connected quartic plane graphs, simple 3-connected quartic plane graphs, and (simple) 4-regular, 3-connected, 6-cyclically-connected plane graphs. These relationships are all discussed in [17].

It is easy to see that a generation theorem for a class can be translated to a generation theorem for its dual class by taking the planar duals of each starting graph, and of the diagrams of the expansion operations. This also works for radial (and medial) graphs. By this relationship, there is an early generation theorem for the class of graphs in Theorem 3.6(4). Namely, Tutte's wheel theorem [83, 84] for generating 3-connected planar graphs translates in the medial form to the result that 3-connected quadrangulations without separating 4-cycles can be generated from the pseudo-double wheels using E_4, which is a more general version of E_1.
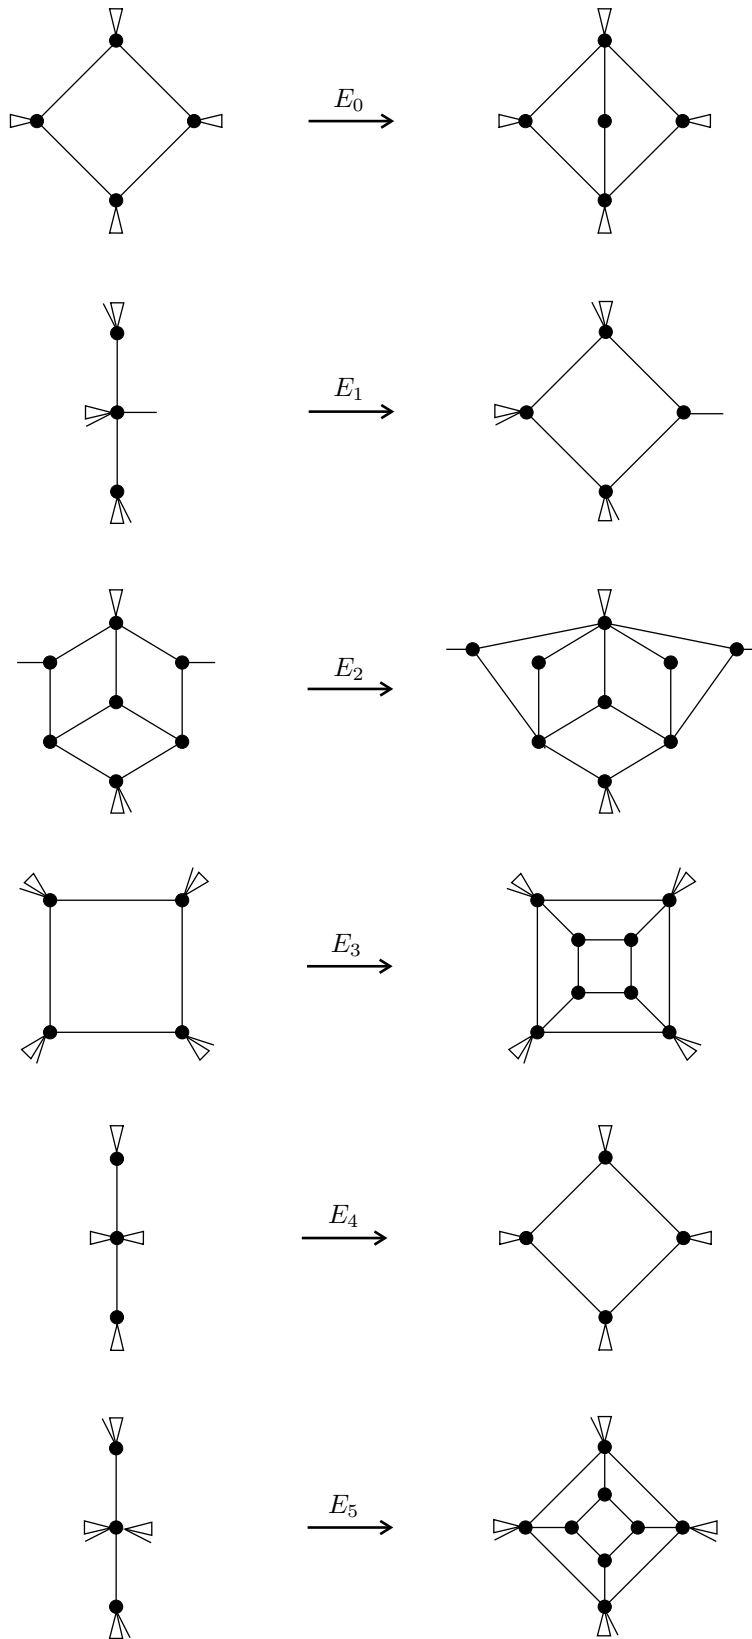
Figure 3.2: Expansion operations for generating simple quadrangulations.

For simple quadrangulations, a recursive structure was given by Negami and Nakamoto [66] and Batagelj [6] starting with the square and using the expansion E_4. There are also several more recent approaches that are geared toward specific theoretical applications. Bau, Matsumoto, Nakamoto and Zheng [8] define a recursive structure on this class which, in terms of reductions, uses E_0-reductions and some operations called hexagonal contractions which both have the property that the reduced graph is a minor of the original. Another approach by Fuchs and Gellert [36] also replaces E_1-reductions by t-contractions, which are operations that come up in the area of t-perfect graphs, and use this to characterise some classes of t-perfect graphs.

The simple quadrangulations with minimum degree 3 was first given a recursive structure by Nakamoto, using the same starting class but with operation E_4 instead of E_1. Again, the updated version is stronger in the sense that E_1 is a more restrictive operation; there are linearly many E_1 expansions possible in a given simple quadrangulation, but E_4 may permit a quadratic number of expansions [17].

A generation theorem for 3-connected quadrangulations that are not necessarily simple was proposed (in the dual form) by Manca [58], however, an error was later found by Lehel [56] who corrected it by adding some expansion operations. Restricting to the subclass of simple graphs, Broersma, Duijvestijn and Göbel [23] gave a generation theorem starting from the cube and using the operations E_3, E_4 and a less restricted version of E_2. More recently, Suzuki [82] gave another alternative to Theorem 3.6(3) in which E_3 is replaced by E_5, which is known as a cube contraction.

The classes of quadrangulations we have mentioned so far all consist of simple graphs. This is representative of results in literature with an exception being that Kápolnai, Domokos, and Szabó [52] give a generation theorem for the class of arbitrary quadrangulations with no restrictions. We will discuss this in more detail when we define our own further classes of quadrangulations.

Pentangulations and beyond

As soon as we move past quadrangulations to k-angulations for $k \geq 5$, the structures turn out to be much more difficult and there are comparatively fewer generation results. This gap was addressed by Hasheminezhad, McKay and Reeves [43] who give recursive structures for simple pentangulations with minimum degree 2, and simple planar 5-regular graphs (these are not dual). Before this, generation

theorems for k-edge-connected 5-regular planar graphs that are not necessarily simple for k up to 5 were also given by Ding, Kanno and Su [28], where the $k = 0$ case refers to arbitrary not necessarily connected graphs. This implies a generation theorems for pentangulations with girth k, so in particular the dual to the $k = 3$ case generates the simple pentangulations.

In Proposition 2.10, we showed that there are no simple planar k-regular graphs for $k > 5$. However, there is no such limit on simple k-angulations. The recursive generation of these classes has been studied by Jooyandeh [49].

3.1.2 Two related classes

For completeness, we will also record some generation theorems for related classes of graphs that are not necessarily planar.

General quartic graphs

Ding and Kanno [27] give generation theorems for quartic graphs that are not necessarily planar or simple. These classes are the arbitrary quartic graphs, connected quartic graphs, loopless quartic graphs, 2-edge-connected loopless quartic graphs, 4-edge-connected loopless quartic graphs, simple quartic graphs, connected simple quartic graphs and 4-edge-connected simple quartic graphs. Before this, a recursive structure for the simple quartic graphs had earlier been given by Bories, Jolivet and Fouquet [12].

k-angulations on surfaces

We have mentioned that graphs can be embedded on surfaces other than the sphere, and this leads to analogous generation theorems for classes of graphs on fixed surfaces. Mirroring the planar case, there is a substantial body of work on generating triangulations on higher genus surfaces, a good deal of work on quadrangulations and, to the best of our knowledge, nothing on higher k-angulations.

The results in this area have a slightly different feel to their planar counterparts. On surfaces, it is often the case that we fix some basic operation, so the new difficulty lies in determining the appropriate starting set of graphs to get a generation theorem. These starting graphs are called *irreducible* with respect to some expansion operation, in reference to the property that they cannot be reduced by the inverse operation into a smaller graph in the class.

For quadrangulations, the most widely considered reduction operation is face contraction, which is operation E_4 shown in Figure 3.2. It is proved in [67] that an E_4-irreducible quadrangulation of a closed surface Σ has at most $186(2 - \chi(\Sigma)) - 64$ vertices, where $\chi(\Sigma)$ denotes the Euler characteristic. This implies that there are finitely many E_4-irreducible quadrangulations (up to equivalence) for any closed surface. So far, without repeating the spherical case which we have already seen, these have been determined completely for the projective plane [66], Klein bottle ([63], corrected in [80]), and torus [65]. In addition, for 3-connected quadrangulations, it is known that E_4-irreducibility is equivalent to E_3-irreducibility [62], and those that are E_5-irreducible have been determined for the projective plane [82].

3.2 More classes of quadrangulations

We now work toward bridging the gap between the results of Brinkmann et al. [17] on simple quadrangulations, Kápolnai, Domokos, and Szabó [52] on the class of arbitrary multiquadrangulations, and those of Ding and Kanno [27] on quartic graphs that are not necessarily planar and not necessarily simple.

The classes of quadrangulations that we consider look, at first glance, a little unusual. We will see that one of them, namely $\mathcal{C}_3$, is the class of 3-edge-connected quadrangulations (Corollary 3.17). Secretly though, the underlying goal is to obtain recursive structures for a sequence of nice classes of quartic plane graphs that are similar to the results we have seen; they are defined by imposing certain conditions on connectedness and allowable non-simple structures. To phrase this in the more convenient land of quadrangulations, we need to characterise the duals of the classes we wish to generate.

The following characterisation is new, and will given, on the quadrangulation side, in terms of the forbidden families of graphs defined via the diagrams in Figure 3.4. When we refer to the *defining* edges, vertices and faces of these structures, we mean the ones that are explicitly shown in this figure. We say that a graph is a member of a specific family if it is represented by the familial diagram. In general, this representation will not be unique; it is often possible for a graph to be depicted as shown with different choices of defining edges, and a graph may be a member of more than one family. In particular, the family of e-sunnies is contained within the family of sunnies.
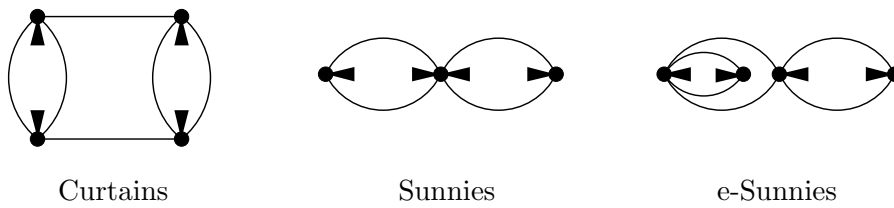
Figure 3.3: Forbidden families in some classes of quadrangulations.

Let δ denote a lower bound on the minimum degree. A standard application of Euler's formula shows that every quadrangulation has average degree strictly less than 4 and hence has minimum degree at most 3. This leads to the following classes of quadrangulations, where $\delta = 2$ or 3:

$\mathcal{A}_\delta$ Simple quadrangulations of the sphere.

$\mathcal{B}_\delta$ Quadrangulations of the sphere excluding curtains and sunnies.

$\mathcal{C}_\delta$ Quadrangulations of the sphere excluding curtains and e-sunnies.

$\mathcal{D}_\delta$ Quadrangulations of the sphere excluding sunnies.

$\mathcal{E}_\delta$ Arbitrary quadrangulations of the sphere with given minimum degree.

$\mathcal{E}_1$ Arbitrary quadrangulations of the sphere.
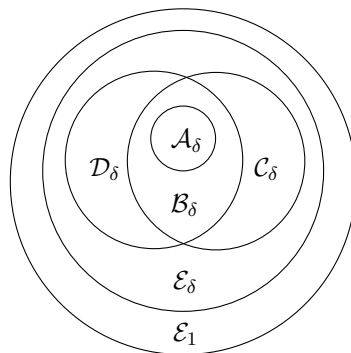
Figure 3.4: Inclusions of our classes of quadrangulations, fixing δ as 1 or 2.
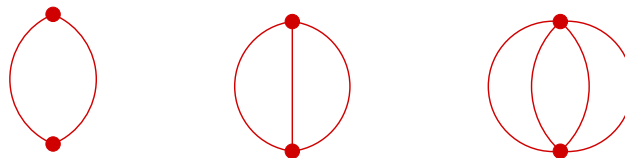
Figure 3.5: All possible clusters of facial digons in the families of quartic graphs. The rightmost one only occurs as the dual graph of the square (C_4).

Here's a handy list of the classes of quartic plane graphs dual to each of the classes defined above. A cluster of facial digons refers to the structures shown in Figure 3.5. We defer the proofs until Section 3.2.2.

$\mathcal{A}_3^*$ Simple 4-edge-connected quartic plane graphs. (Theorem 3.14)

$\mathcal{A}_2^*$ 4-edge-connected quartic plane graphs. (Theorem 3.14)

$\mathcal{B}_3^*$ Simple 2-connected quartic plane graphs. (Theorem 3.15)

$\mathcal{B}_2^*$ 2-connected quartic plane graphs that are simple except for those clusters of facial digons shown in Figure 3.5. (Theorem 3.15)

$\mathcal{C}_3^*$ Simple quartic plane graphs. (Theorem 3.16)

$\mathcal{C}_2^*$ Quartic plane graphs that are simple except for clusters of facial digons. (Theorem 3.16)

$\mathcal{D}_3^*$ 2-connected quartic plane graphs without facial digons or loops. (Theorem 3.18)

$\mathcal{D}_2^*$ Loopless 2-connected quartic plane graphs. (Theorem 3.18)

$\mathcal{E}_3^*$ Quartic plane graphs without facial digons or loops. (Theorem 3.19)

$\mathcal{E}_2^*$ Loopless quartic plane graphs. (Theorem 3.19)

$\mathcal{E}_1^*$ Arbitrary quartic plane graphs. (Theorem 3.19)

3.2.1   Generation theorems

The operations used to generate a particular class of quadrangulations from an earlier class are known as expansions. Applying an expansion will not always result in a strictly larger quadrangulation, but it will introduce some more general structure. We will use the five expansion operations shown in Figure 3.6. To make the displayed expansion operations precise, we detail their associated reductions as follows.

- Suppose we have a 6-cycle $v_1 \overset{e_1}{\text{---}} v_2 \overset{e_2}{\text{---}} v_3 \overset{e_3}{\text{---}} v_4 \overset{e_4}{\text{---}} v_5 \overset{e_5}{\text{---}} v_6 \overset{e_6}{\text{---}} v_1$ together with a digon $v_1 \overset{c_1}{\text{---}} v_4 \overset{c_2}{\text{---}} v_1$ such that $v_1 \overset{e_1}{\text{---}} v_2 \overset{e_2}{\text{---}} v_3 \overset{e_3}{\text{---}} v_4 \overset{c_1}{\text{---}} v_1$ and $v_1 \overset{c_1}{\text{---}} v_4 \overset{e_4}{\text{---}} v_5 \overset{e_5}{\text{---}} v_6 \overset{e_6}{\text{---}} v_1$ are facial cycles. A P-reduction consists of deleting the c_2 and inserting an edge between v_3 and v_6. By planarity, there is only one way to do this up to equivalence so the cyclic order at each vertex is completely determined.

- A P_2-reduction is the same as a P-reduction except that it can only be applied when e_2 and e_5 are simple.

- A P_1-reduction is also the same as a P-reduction except that it can only be applied when e_2 is simple, there exists an edge e_5' parallel to e_5 which does not share a face with c_2, and the region bounded by $v_1 \overset{e_6}{-\!\!-} v_6 \overset{e_5'}{-\!\!-} v_5 \overset{e_4}{-\!\!-} v_4 \overset{c_1}{-\!\!-} v_1$ is non-facial.

- A Q_2-reduction may be applied when there is a face bounded by a circuit of the form $v_1 \overset{e_1}{-\!\!-} v_2 \overset{e_2}{-\!\!-} v_1 \overset{e_3}{-\!\!-} v_3 \overset{e_4}{-\!\!-} v_1$, with the condition that v_1 and v_3 are both incident to edges inside the region bounded by the digon $v_1 \overset{e_4}{-\!\!-} v_3 \overset{e_3}{-\!\!-} v_1$ that does not contain v_2. The reduction then consists of identifying v_2 and v_3 to form a new vertex v_2', as well as identifying e_1 with e_4 and e_2 with e_3 to form new edges e_1' and e_2' that both have as endvertices v_1 and v_2'.

- A Q_1-reduction may be applied when there is a face bounded by a walk of the form $v_1 \overset{e_1}{-\!\!-} v_2 \overset{e_2}{-\!\!-} v_1 \overset{e_3}{-\!\!-} v_3 \overset{e_3}{-\!\!-} v_1$, and consists of identifying v_2 and v_3 to form a new vertex v_2', as well as identifying all three edges e_1, e_2 and e_3 into a single edge e_1' with endvertices v_1 and v_2'.

We can already generate $\mathcal{A}_\delta$ (Theorem 3.6(1) and (2)), so we will use these are starting classes to find recursive structures for each of the later classes in sequence. Kápolnai, Domokos, and Szabó [52] showed that $\mathcal{E}_1$ can be generated from P_2 using the operations E_0, E_1 and Q_1. Our generation theorem for this class is not as efficient as theirs since we use more operations. Nonetheless, the scenic route that we take allows us to obtain generation theorems for the other eight new subclasses of quadrangulations.

For each of the following generation theorems, the idea is to show that any graph in the target class is reducible via the specified operation(s), and that the reduced graph in each case is still in the class. This would then imply that the required sequence of graphs can be found by recursively applying reductions. The following face lemma will help us to do this.

Lemma 3.7. *All faces of a quadrangulation of the sphere take one of the types A, B, C or D shown in Figure 3.7.*

Proof. Each face is bounded by some 4-cycle $v_1 \overset{e_1}{-\!\!-} v_2 \overset{e_2}{-\!\!-} v_3 \overset{e_3}{-\!\!-} v_4 \overset{e_4}{-\!\!-} v_1$. To enumerate the possible faces, we need to consider all possible ways that some of these edges and vertices might be the same. Firstly, if all edges and vertices are distinct then we have face type A. Otherwise we must have at least two vertices be the same. Since there are no loops, it is only possible that $v_1 = v_3$ or $v_2 = v_4$. Moreover, it is clear that at most one of these can hold if we are to end up with one face. Without
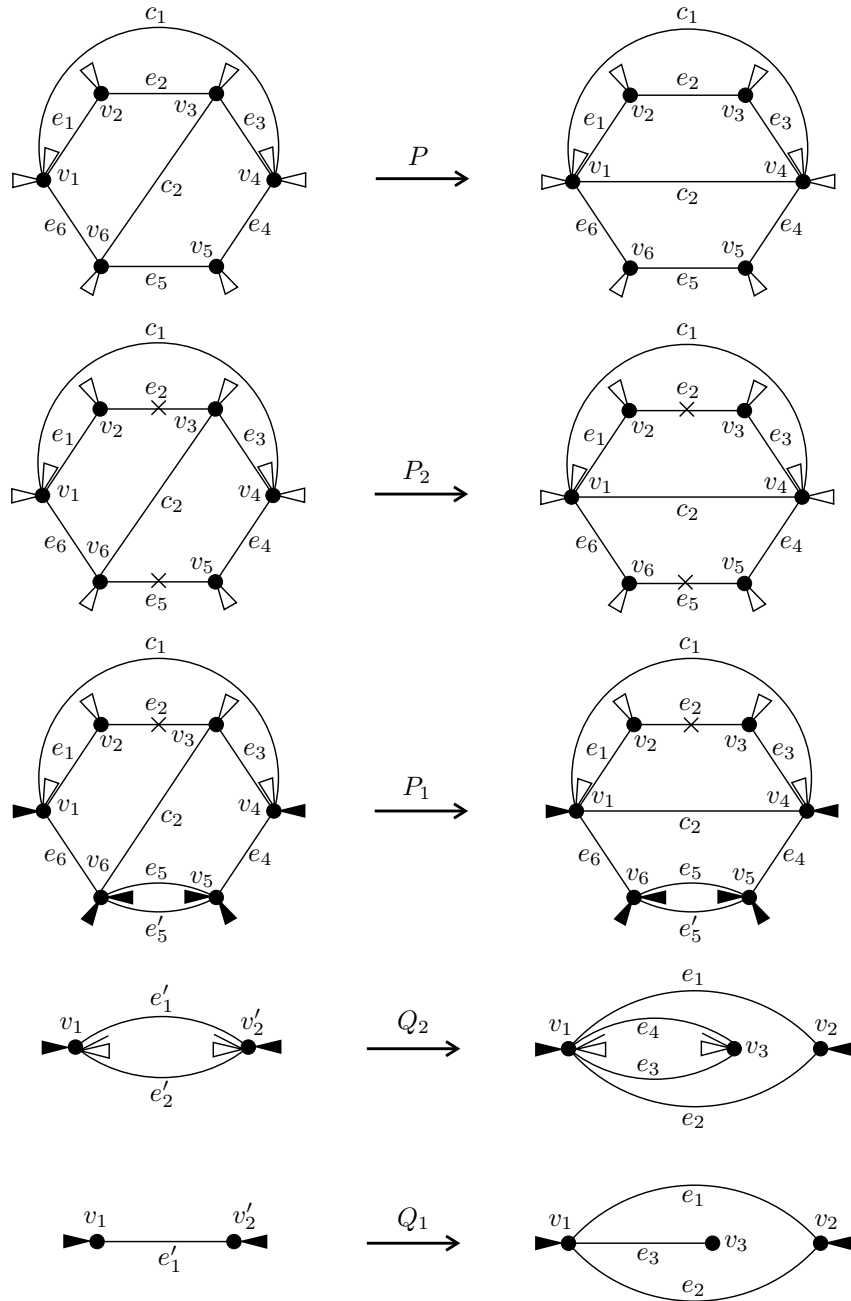
Figure 3.6: Expansion operations. Edges marked with a cross are required to be simple.

loss of generality, suppose $v_2 = v_4$. If all edges are then distinct, we obtain a face
of type B, shown in the figure as both a bounded region or as an outside face.
Otherwise, the only possible edge identifications are $e_2 = e_3$ or $e_1 = e_4$, the other
possibilities being excluded since they would create loops. When exactly one of
these is true we have a face of type C, again drawn as both an inside and outside
face in the figure, and when both hold we get a face of type D.   □
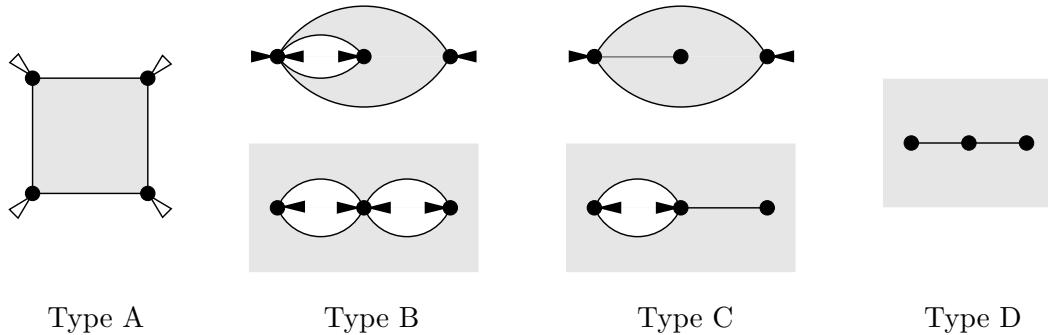

| Type A | Type B | Type C | Type D |

Figure 3.7: All possible faces of length four.

We are now ready to state our first generation theorem in the sequence.

Theorem 3.8. *The class $\mathcal{B}_\delta$ ($\delta = 2, 3$) of all quadrangulations of the sphere
excluding curtains and sunnies is generated from $\mathcal{A}_\delta$ by the operations P_2 and P_1.*

Proof. Let G be a graph in $\mathcal{B}_\delta - \mathcal{A}_\delta$. We first show that every face in G is
bounded by a simple cycle. This follows from Lemma 3.7 since of the possible
faces in Figure 3.7, types C and D cannot occur as $\delta \geq 2$ and any graph containing
a type B face is in the family of sunnies. This leaves only face type A, which is
the one we wanted. Now since $G \notin \mathcal{A}_\delta$, it must be non-simple. Moreover, since
it is bipartite and hence has no loops, this means that it must contain parallel
edges.

Fix a planar embedding of G and pick a pair of innermost parallel edges
c_1 and c_2 with endvertices v_1 and v_4; by innermost we mean the union of the
subgraph contained inside the digon $c_1 c_2$ together with the digon itself is simple.
Consider the simple cycles bounding the two faces adjacent to c_2. Applying the
Jordan curve theorem to the digon $c_1 c_2$, we see that the vertex sets of these two
simple cycles intersect only in v_1 and v_3, and their disjoint union is a 6-cycle
$v_1 \overset{e_1}{\longrightarrow} v_2 \overset{e_2}{\longrightarrow} v_3 \overset{e_3}{\longrightarrow} v_4 \overset{e_4}{\longrightarrow} v_5 \overset{e_5}{\longrightarrow} v_6 \overset{e_6}{\longrightarrow} v_1$. Let's suppose these are labelled so that v_2 and v_3
are inside $c_1 c_2$.

Now consider the edges e_2 and e_5. By choice of c_1 and c_2, we know that e_2
is necessarily simple. If e_5 is also simple then we can apply an P_2-reduction.

On the other hand, if e_5 is not simple then we claim that a P_1-reduction can be applied. Say e_5 is parallel to e_5', labelled so that e_5 and c_2 share a face. The only possible obstruction to the reduction is if $v_1 \overset{e_6}{\rule{0.8em}{0.4pt}} v_6 \overset{e_5'}{\rule{0.8em}{0.4pt}} v_5 \overset{e_4}{\rule{0.8em}{0.4pt}} v_4 \overset{c_1}{\rule{0.8em}{0.4pt}} v_1$ is a facial cycle. However, we observe that if this is the case then the graph is a curtain as shown in Figure 3.8, which cannot occur.
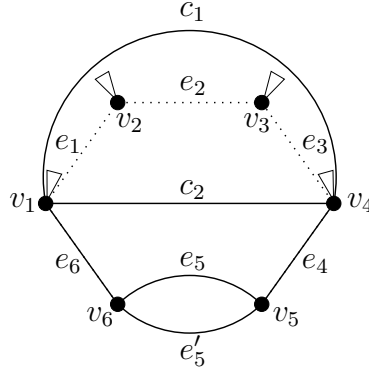

Figure 3.8: The solid lines form a forbidden curtain.

It remains to show that after applying either of these reductions, the resulting graph is still in B_δ. Since both operations remove a parallel edge and cannot create a new one due to the ordering of c_1 at v_1 and v_4, it is not possible for either reduction to produce a graph in the family of curtains or sunnies. For the degree condition, it is enough to note that $d(v_1)$ and $d(v_4)$ are at least 4 prior to the reduction, and so are at least 3 afterward. No other vertices decrease in degree. $\qquad\square$

Theorem 3.9. *The class $\mathcal{D}_\delta$ of all quadrangulations of the sphere excluding sunnies is generated from $\mathcal{B}_\delta$ by operation P.*

Proof. Take any G in $\mathcal{D}_\delta - \mathcal{B}_\delta$. The same argument as before shows that every face in G is bounded by a simple cycle. From the definitions of the classes, G must be a curtain, so in particular let $v_1 \overset{c_1}{\rule{0.8em}{0.4pt}} v_4 \overset{c_2}{\rule{0.8em}{0.4pt}} v_1$ be one of the defining digons. We can now follow the previous proof, but where we differentiated between the two operations previously, we find that in all cases the graph is P-reducible. The verification that degree bounds are preserved and the reduced graph is not in the family of sunnies still holds. $\qquad\square$

Theorem 3.10. *The class $\mathcal{C}_\delta$ of all quadrangulations of the sphere excluding curtains and e-sunnies is generated from $\mathcal{B}_\delta$ by operation Q_2.*

Proof. If G is in $\mathcal{C}_\delta - \mathcal{B}_\delta$, then it must contain sunnies that are not e-sunnies. To start with, we will just use the existence of type B faces. Pick any type B face, say bounded by the circuit $v_1 \overset{e_1}{-\!\!\!-} v_2 \overset{e_2}{-\!\!\!-} v_1 \overset{e_3}{-\!\!\!-} v_3 \overset{e_4}{-\!\!\!-} v_1$, and fix an embedding of G so that this face is not the outer face. Now consider the digon $v_1 \overset{e_4}{-\!\!\!-} v_3 \overset{e_3}{-\!\!\!-} v_1$. This is necessarily non-facial, so let H denote the subgraph of G contained inside this digon. Suppose H is adjacent to only one of v_1 or v_3. If it is adjacent to v_3, then there must be a vertex v_4 and parallel edges e_5 and e_6 between v_3 and v_4 in order to ensure the face inside the digon adjacent to e_3 and e_4 has length four. Furthermore, e_5 and e_6 must be distinct, otherwise v_4 would have degree 1. But G is not in the family of e-sunnies, so H must be adjacent to v_1 only. At this point we have $d(v_3) = 2$, so if $\delta = 3$ then this situation is impossible. Thus, H must be adjacent to both v_1 and v_3 and we may apply an Q_2-reduction. Note that since $e_1 e_2$ has length 2, the region outside this digon is non-facial since G is a quadrangulation.

For $\delta = 2$, we can still show that there is exists a type B face for which H is adjacent to both v_1 and v_3. Under the previous labelling, if have H adjacent to only v_1 then as before there is a face inside the digon $v_1 \overset{e_3}{-\!\!\!-} v_3 \overset{e_4}{-\!\!\!-} v_1$ that is adjacent to both e_3 and e_4. In order for this to be bounded by a walk of length 4, we must be now have a vertex v_4 and distinct parallel edges e_5 and e_6 between v_1 and v_4 so that $v_1 \overset{e_4}{-\!\!\!-} v_3 \overset{e_3}{-\!\!\!-} v_1 \overset{e_5}{-\!\!\!-} v_4 \overset{e_6}{-\!\!\!-} v_1$ is a face of G. This is gives another type B face, as can be seen in Figure 3.9, which we shall declare to be inside our original one. We now have the notion of an innermost type B face. Choosing such a face and again using the same labelling, the property that it is innermost guarantees that H must be adjacent to both v_1 and v_3 so as before we find that G is Q_2-reducible.
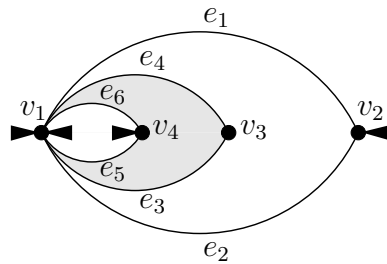

Figure 3.9: An inner type B face (shaded) inside a larger one.

Suppose that the reduced graph is in the family of curtains. Then $e_1' e_2'$ must be a defining digon in a curtain structure since the rest of the graph is unchanged, and the other defining edges are either all inside or all outside $e_1' e_2'$. In either case, we see that these edges are still present before the reduction, and together

with either e_1e_2 or e_3e_4 allow us to draw G as a member of the curtain family which is a contradiction. Similarly, the reduced graph cannot be in the family of e-sunnies. Also, before the reduction $d(v_1) \geq 5$ so the degree is at least 3 afterward, and no other vertices that survive the reduction have lower degrees. Therefore, $Q_2(G) \in \mathcal{C}_\delta$. $\qquad \square$

Theorem 3.11. *The class $\mathcal{E}_\delta$ of all quadrangulations of the sphere with given δ is generated from $\mathcal{D}_\delta$ by operation Q_2.*

Proof. We use similar reasoning as in the proof of Theorem 3.10. If $G \in \mathcal{E}_\delta - \mathcal{D}_\delta$, then it must be in the family of sunnies, and consequently has a type B face. Using the notation and argument in the $\delta = 2$ case of the previous proof, one can find an innermost type B face on which to apply a Q_2-reduction. The argument that the degree bounds are preserved also carries over. $\qquad \square$

Theorem 3.12. *Except for the path of three vertices, the class $\mathcal{E}_1$ of arbitrary quadrangulations of the sphere is generated from $\mathcal{E}_2$ by operation Q_1.*

Proof. Let G be a graph in $\mathcal{E}_1 - \mathcal{E}_2$, meaning it must have a degree 1 vertex. By Lemma 3.7, any face containing a vertex of degree 1 is of type C or type D. The latter is unique to the case where G is a path with three vertices, so in all other graphs we must have a type C face. This is exactly the form needed to apply a Q_1-reduction, and doing so cannot produce isolated vertices so the degree condition is met. $\qquad \square$

3.2.2 Characterising the duals

The proofs of the dual relationships we listed earlier are relatively straightforward once we attain some fluency in translating between various properties of a graph and the corresponding properties or structures of its dual. We have already seen cut-cycle duality (Lemma 2.23), which gives one result of this type. The following lemma spells out the correspondence between key structures that appear in our classes and their duals.

Lemma 3.13 (Dual structures)**.** *Suppose G is a quadrangulation, or equivalently that G^* is quartic.*

(a) *If G is simple and G^* is 4-edge-connected, then G has a vertex of degree 2 if and only if G^* has parallel edges.*

(b) *G is in the family of sunnies if and only if G^* has a 1-vertex-cut.*

(c) G is in the family of curtains or e-sunnies if and only if G^* has a non-facial digon.

(d) If G^* is 2-vertex-connected, then G is in the family of curtains if and only if G^* has a non-facial digon.

(e) G has a degree 2 vertex if and only if G^* has a facial digon.

(f) G has a degree 1 vertex if and only if G^* has a loop.

Proof. (a) By examining the picture on the right of Figure 3.10, one observes that if G has a vertex of degree 2 (this is a 2-edge cut), then the dual, drawn with unfilled vertices and red edges, must have a digon so this direction is straightforward. Suppose instead that we start with a digon in G^*. If the region inside the digon were a face, then the vertex corresponding to this face in the dual would have degree two, so we would be done in this case. Otherwise, there is some part of the graph G^* contained in the shaded region as shown in the left of Figure 3.10. Since G^* is 4-edge-connected, there must be four edges connecting this part of the graph to the digon, and the configuration drawn with two edges between the shaded region and each vertex of the digon is the only possibility. However, we have now accounted for all four edges adjacent to the vertices in the digon, and since G^* is connected, this means that all of G^* must actually be within the shaded region. In particular, the outside region of the digon is a face, and hence corresponds to a vertex of degree 2 in G.
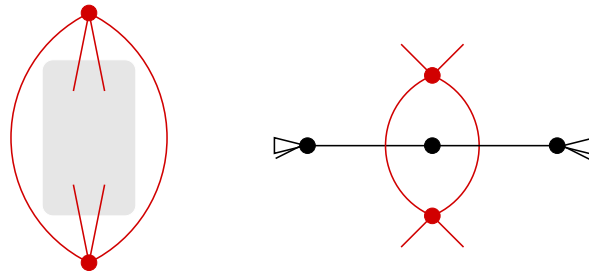

Figure 3.10: Degree two vertices in $\mathcal{A}_3$ are dual to facial digons in $\mathcal{Q}$.

(b) Suppose $\{v\}$ is a 1-vertex-cut in G^*. Using the unique form of a 2-cut discussed in Section 2.3.1, we must be able to draw G in the form shown in Figure 3.11 where all vertices and edges not drawn are contained in one of the shaded regions. The dual of the graph is shown with unfilled vertices and red edges, and demonstrates that G is in the family of sunnies since the outside face is of type B (refer to Figure 3.7). Conversely, if G^* has a type B face then we

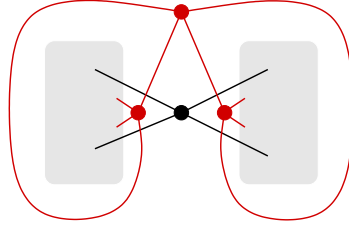see from the same figure that G has a 1-cut consisting of the vertex dual to the outer face.

Figure 3.11: Sunnies are dual to 1-cuts

(c) Suppose there is a non-facial digon in G^*. Since it is quartic, there are four possible configurations of edges at the two vertices in the digon, and these are shown in Figure 3.12. From the left, the 00-digon is facial so we will not consider it. The 01-configuration has an unbalanced 1-vertex-cut which violates our characterisation in Section 2.3.1. The last two configurations are possible.
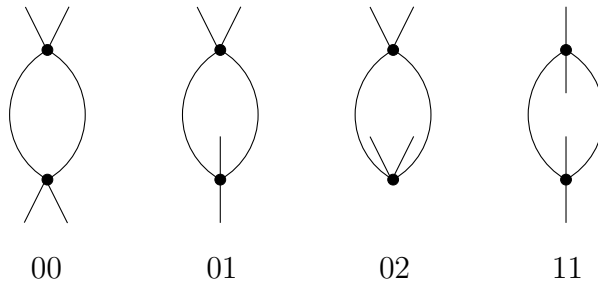

$$00 \qquad 01 \qquad 02 \qquad 11$$

Figure 3.12: Possible configurations of digons in a quartic graph.

If the 02-configuration is present, the graph has the structure of the red graph on the right of Figure 3.13. In the dual G, drawn in solid black, observe that in each region bounded by the digon the two vertices of G shown must exist and be distinct since G being loopless implies that G^* cannot have any 1-edge-cuts. Thus, G is in the family of e-sunnies.
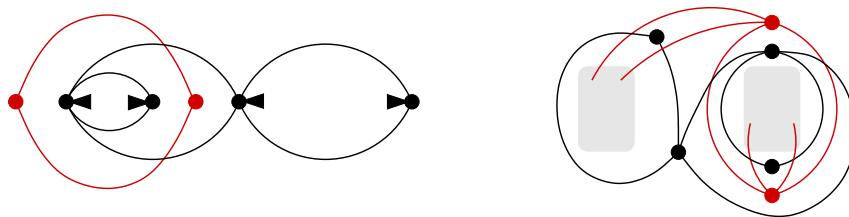

Figure 3.13: e-Sunnies are dual to 02-digons.

Similarly, if the 11-configuration is present then G^* has the structure of the red graph on the right of Figure 3.14. Again, the shaded regions must be connected subgraphs of G^* to avoid 1-edge-cuts, and we see that the dual G is in the family of curtains.
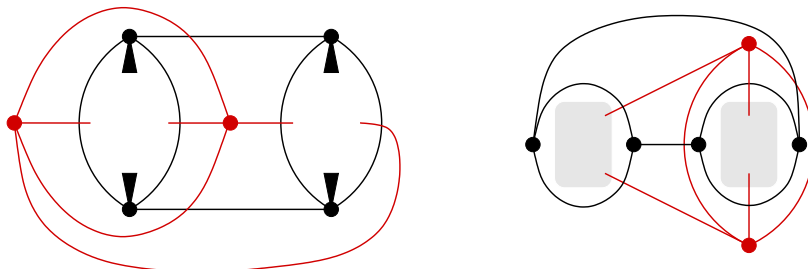

Figure 3.14: Curtains are dual to 11-digons.

Conversely, if G is in the family of e-sunnies or curtains then the pictures on the left hand side of Figure 3.13 and Figure 3.14 respectively show that we get either a 11- or 02-digon in G^*, both of which are non-facial.

(d) If G is in the family of curtains then we have already seen in (c) that G^* has a digon, and moreover by inspecting Figure 3.14 we observe that it is non-facial. The converse can also be shown easily following the argument from part (c). Returning to the digon configurations from Figure 3.12, we may still ignore the first two possibilities. The assumption that G^* is 2-connected now also excludes the 02-configuration as both vertices in the digon are 1-cuts. Thus, only the fourth configuration can occur, and we have already shown that this corresponds G being in the family of curtains.

(e) This is a basic observation.

(f) Certainly if G has a degree 1 vertex, then the dual has a loop. Conversely, if G^* has a loop then G must have a bridge, since removing the edge dual to the loop will disconnect the graph. We observe that any bridge in a quadrangulation is adjacent to a vertex of degree 1 since there are no loops and the bridge must occur twice in the facial circuit that contains it.   □

The previous lemma contains almost all the work needed to determine the dual classes of each of our classes of quadrangulations. What remains amounts to some piecing together of those structures mentioned to obtain the defining properties of the classes.

Theorem 3.14. $\mathcal{A}_3^*$ *is the class of simple 4-edge-connected quartic plane graphs, and* $\mathcal{A}_2^*$ *is the class of 4-edge-connected quartic plane graphs.*

Proof. Let Q4E denote the class of 4-edge-connected quartic plane graphs, and QS4E the class of simple 4-edge-connected quartic plane graphs. We begin with the second statement. It is elementary that quadrangulations are dual to quartic plane graphs, so we only need to check the remaining conditions. Let G be any graph in $\mathcal{A}_2$. Observe that G has no loops or digons since it is simple, and no 3-cycles since it is bipartite so by cut-cycle duality it follows that G^* is 4-edge-connected. Thus, $\mathcal{A}_2^* \subseteq$ Q4E. Conversely, if H is some graph in Q4E, then H is 4-edge-connected and cut-cycle duality immediately implies that H^* has no loops or parallel edges. Any graph is isomorphic to its double dual, so this shows that Q4E $\subseteq \mathcal{A}_2^*$ which now gives the second statement.

Since $\mathcal{A}_3^* \subset \mathcal{A}_2^*$, we already know that the dual of any graph $G \in \mathcal{A}_3$ is a 4-edge-connected plane quartic graph. To prove that $\mathcal{A}_3^* =$ QS4E now, we just need to show that a graph in Q4E is simple if and only if its dual has minimum degree three. Since simple quadrangulations are bridgeless and hence their duals are loopless, then this is equivalent to showing that a graph in $\mathcal{A}_2^*$ has a vertex of degree two if and only if its dual has parallel edges which is the content of Lemma 3.13(a). □

Theorem 3.15. $\mathcal{B}_3^*$ *is the class of simple 2-connected quartic plane graphs, and* $\mathcal{B}_2^*$ *is the class of 2-connected quartic plane graphs that are simple except for facial digons.*

Proof. Let QS2 be the class of simple 2-connected quartic planar graphs, and QS2$^+$ denote the class of 2-connected quartic planar graphs that are simple except for facial digons. We will start by showing the first part of the statement. Suppose $G \in \mathcal{B}_3$. It is straightforward to establish that G^* is a loopless quartic plane graph. By Lemma 3.13(b), since sunnies are forbidden from $\mathcal{B}_3$ it follows that G^* must be 2-connected. As G has no curtains, G^* has no non-facial digons by Lemma 3.13(d), and as $\delta = 3$ it follows from Lemma 3.13(e) that G^* has no facial digons either. That is, G^* has no parallel edges and hence is simple. Therefore, $\mathcal{B}_3^* \subseteq$ QS2. Starting instead with an arbitrary graph $H = G^* \in$ QS2, then G is a quadrangulation. It has no degree 1 or 2 vertices by Lemma 3.13(e) and (f) since G is simple. In addition, parts (b) and (d) of Lemma 3.13 ensure that G is not in the family of sunnies or curtains respectively, using the properties of G being 2-connected and G^* having no parallel edges (specifically, no non-facial digons).

Simply omitting those statements in the argument above that require $\delta = 3$ gives the second statement. Explicitly, we have without any modification to the proof that the dual of $G \in \mathcal{B}_2$ is 2-connected, quartic and without loops or non-facial digons which gives one inclusion of the second statement. Conversely, if $G^* \in \text{QS2}^+$ then the above arguments show that G is a loopless quadrangulation not in the family of sunnies or curtains. Hence, $\mathcal{B}_2^* = \text{QS2}^+$.   □

Theorem 3.16. *$\mathcal{C}_3^*$ is the class of simple quartic plane graphs, and $\mathcal{C}_2^*$ is the class of quartic plane graphs that are simple except for facial digons.*

Proof. Let QS be the class of simple quartic plane graphs, and QS^+ denote the class of quartic plane graphs that are simple except for facial digons. If $G \in \mathcal{C'}_3^*$, then G^* is quartic plane graph with no loops. Lemma 3.13(c) and (d) guarantee that it also has no digons, so G^* simple. Therefore, $G^* \in \text{QS}$ and $\mathcal{C'}_3^* \subseteq \text{QS}$. For $G^* \in \text{QS}$, the same lemmas show that G is not in the family of curtains or e-sunnies. From Lemma 3.13(e) and (f), G cannot have any vertices with degree less than 3, so $\mathcal{C'}_3^* = \text{QS}$.

As in the last proof, to obtain the second statement with $\delta = 2$ we may simply repeat the previous argument omitting those statements that require the degree bound.   □

Corollary 3.17. *$\mathcal{C}_3$ is the class of 3-edge-connected quadrangulations of the sphere.*

Proof. By Theorem 3.16, we have $\mathcal{C'}_3 \cong \mathcal{C'}_3^{**} \cong \text{QS}^*$ which, using cut-cycle duality, is the class of 3-edge-connected quadrangulations of the sphere.   □

Theorem 3.18. *$\mathcal{D}_3^*$ is the class of 2-connected quartic plane graphs without facial digons or loops, and $\mathcal{D}_2^*$ is the class of 2-connected quartic plane graphs without loops.*

Proof. By comparison with Theorem 3.15, it is enough to show that a graph $G \in \mathcal{D}_3^*$ has curtains if and only if G^* has a non-facial digon. Given 2-connectedness of G^* from Lemma 3.13(b), this is precisely Lemma 3.13(d).   □

Theorem 3.19. *$\mathcal{E}_3^*$ is the class of quartic plane graphs without facial digons or loops, $\mathcal{E}_2^*$ is the class of quartic plane graphs without loops, and $\mathcal{E}_1^*$ is the class of arbitrary quartic plane graphs.*

Proof. Lemma 3.13(e) and (f) allow us to deduce the first two statements. The final claim is clear.   □

3.2.3 Translating to quartic planar graphs

The results of the previous two sections, together with the bijections that we get from taking planar duals, also imply that we have generation theorems for each of our nine dual classes of quartic plane graphs. If we follow through the chain of theorems back to those for the simple quadrangulations then the basic starting graphs are the duals of the pseudo-double wheels, which are isomorphic to the antiprisms shown in Figure 3.15, or the dual of the square.
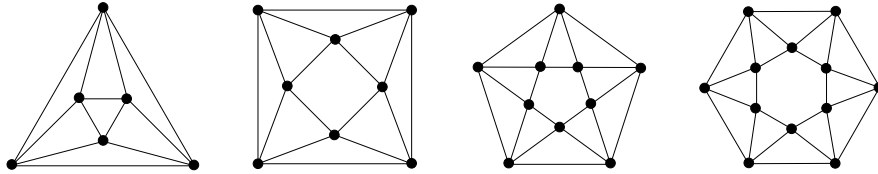

Figure 3.15: The n-antiprisms for $n = 3, 4, 5, 6$.

To illustrate this, we state the result for the class of simple quartic plane graphs without explicitly writing down the expansion operations.

Theorem 3.20. *The class $\mathcal{C}_3^*$ of simple quartic plane graphs can be generated from the antiprisms by the operations E_1^*, E_3^*, P_1^*, P_2^* and Q_2^*.*

3.3 Isomorph-free graph generation

We have already mentioned that a primary reason to work on recursive structures of graphs is for the implementation of practical graph generators. Generation theorems guarantee that if all expansion operations are applied in all possible ways, then we will end up with an exhaustive list of graphs in the family. However, this will generally produce many copies of isomorphic graphs which is problematic since we usually work with graphs up to isomorphism. The aim of a computer program should therefore be to list exactly one member of each isomorphism class of graphs in the family, necessitating a method of isomorph rejection.

In 1998, McKay [59] introduced the *canonical construction path* (CCP) which avoids explicit isomorphism testing. This is highly desirable as graph isomorphism is difficult, and naive approaches of isomorph rejection, such as checking graphs pairwise, is very limited in its applicability by memory constraints. Essentially, given a class $\mathcal{C}$ of graphs and a generation theorem, the CCP method builds a unique sequence of expansions that can be used to obtain each graph. This requires the definition of a canonical reduction for each graph. Then, whenever

an expansion is applied to G_i giving a graph G_{i+1}, we check if the inverse of the expansion is the canonical reduction for G_{i+1}. If not, then G_{i+1} is rejected, and those graphs that are not rejected comprise one graph in each isomorphism class. There is much more machinery at work than is alluded to in this rough sketch. For the full implementation, we refer to [22].

The CCP method has been used, for example, in the program `plantri` due to Brinkmann and McKay [20, 22] that generates many classes of planar graphs including planar triangulations (as its name suggests) and quadrangulations, including all of those mentioned in Theorem 3.6. The resulting counts for these graphs are recorded in [17]. The table in Figure 3.16 is extracted from that paper and gives the number $q(n)$ of simple 3-connected quadrangulations on n vertices (and hence with $f = n - 2$ faces) up to possibly reflectional isomorphism. The original table goes up to 36 vertices, on which there are just over three trillion such quadrangulations. In the dual, these are also the counts of 3-connected quartic planar graphs on f vertices.

n	f	$q(n)$
8	6	1
9	7	0
10	8	1
11	9	1
12	10	3

Figure 3.16: Counts of simple 3-connected quadrangulations on n vertices, equivalently 3-connected quartic planar graphs on f vertices.

There is another program, `surftri`, written by Sulanke that is based on `plantri` and which generates triangulations on surfaces. Details together with corresponding counts for the torus, double torus, and non-orientable surfaces up to genus 12 can be found in [79].

3.4 Further work

We have come up with a sequence of generation theorems which each build upon the last class. Whilst this has enabled us to obtain recursive structures for multiple subclasses of quartic planar graphs in one fell swoop, it is quite possible that some of the later classes can be generated using fewer operations. We already know that this is the case, for instance, with arbitrary quadrangulations.

Away from quartic planar graphs and quadrangulations of the sphere, we see two broad directions in which to take generating theorems for maps. The first is to look to generate restricted subclasses of pentangulations and 5-regular plane graphs – there are good results for classes of cubic and quartic graphs that satisfy certain connectedness or degree conditions, but still very few for the 5-regular case. This is likely to be difficult, however, given the complexity of the operations required in [43].

The second possibility is to work on higher-genus surfaces. The proofs of the theorems we mentioned before for the characterisation of irreducible graphs were all particular to the given surface, and so do not generalise well. This was also the case for early characterisations of irreducible triangulations, until Sulanke considered genus-increasing expansion operations that relate the irreducible triangulations on different surfaces [78, 81]. There are two basic operations: given a triangle, the equivalent to adding a handle to the surface is given by removing two faces and then identifying their boundaries, and similarly, a crosscap can be added by removing a vertex of degree six and identifying each neighbour to its opposite. In both cases, if the faces and vertices involved are sufficiently far apart, then this produdces a triangulation of the original surface with a handle or crosscap added.

The generation theorem implied by this has been implemented to determine complete lists of irreducible triangulations on orientable surfaces up to genus 2 and non-orientable surfaces up to genus 4. One can then generate all quadrangulations of these surfaces by some edges in all possible ways. On the other hand, by defining analogous handle- and crosscap-addition expansions using quadrangular faces, we believe it may be possible to also adapt Sulanke's strategy to generate quadrangulations directly.

Chapter 4

Self-avoiding Eulerian circuits

An Eulerian circuit of a graph is a circuit in which every edge of the graph is traversed exactly once. Any graph that admits such a circuit is said to be *Eulerian*. It is a standard fact, attributed to Euler*, that a simple graph is Eulerian if and only if it is connected and all of its vertices have even degree. If one is asked to find an Eulerian circuit given some Eulerian graph, one can breathe a sigh of relief because this is very easy to do. Here are two well-known algorithms.

1. (Fleury) Starting at an arbitrary vertex, choose a sequence of unused edges such that at each step the subgraph induced by the set of unused vertices is connected. It is straightforward to prove that the resulting circuit uses all edges, with the basic principle being that even degrees ensure that when the circuit enters any vertex except the initial one, there is always an edge available for it to exit.

2. (Hierholzer) Again starting at an arbitrary vertex, choose any sequence of available edges until we form a circuit C. If there are any edges of the graph that are not in the circuit, then we pick a vertex v on C that is adjacent to an unused edge and choose another sequence of available edges until we return to v, thus creating another circuit D. These circuits can be joined by following one vw-subpath in C, then taking D, and finally taking the other wv-subpath in C to return to the start. This is repeated until all edges have been incorporated.

The latter procedure highlights an interesting feature of Eulerian circuits that

---

*For a fun historical diversion, we recommend [75]. In particular, Euler proved necessity in 1735, but a proof of sufficiency was not published until 1873 by Hierholzer.

is present but somewhat obscured in the first approach; provided our graph has vertices of degree greater than 2, any Eulerian circuit will cross back on itself many times and hence contain subcircuits and subcycles. More precisely, we make the following definition.

Definition 4.1. Let $C = x_1 x_2 \ldots x_n$ be an Eulerian circuit in a simple graph. A *subcycle* of C is a subpath $x_i x_{i+1} \ldots x_j$ of C such that $x_i = x_j$ and the vertices $x_i, x_{i+1}, \ldots, x_{j-1}$ are all distinct.

A subcircuit of an Eulerian circuit is defined by dropping the condition that the internal vertices are distinct, but since any subcircuit contains a subcycle, we find it unnecessary to use both terms. This raises the question of whether we can control the lengths of these subcycles for particular classes of graphs. Provided the girth of the graph is small enough, it is generally straightforward to find an Eulerian circuit with a small subcycle, so we are interested in when it is possible to find one that avoids small subcycles.

From a graph theoretic perspective, the study of these restricted Eulerian circuits is partly motivated by their connection to tree, path and cycle decompositions of graphs surrounding which there are numerous famous open problems. We will defer this discussion to the next section. In addition, Eulerian circuits have found applications in other areas such as in bioinformatics, where the problem of assembling DNA fragments after sequencing can in some instances be reduced to finding an Eulerian circuit of a particular auxiliary graph [70]. With a view to maximising the potential for these applications, it is desirable for our proofs to be constructive.

4.1   Formulation and related work

We first state our problem precisely – there are several equivalent formulations in literature, so we will try to remain consistent throughout. Adopting the language of [55, 68], we introduce the following terminology.

Definition 4.2. An Eulerian circuit is *k-locally self-avoiding* for some positive integer k if it contains no subcycles of length at most k. The *girth* $g(C)$ of an Eulerian circuit C is the minimum positive integer ℓ such that there is a subcycle of length ℓ.

It is immediate that a circuit with girth ℓ is k-locally self-avoiding for all $k < \ell$. The following problem was first posed by Häggkvist [41] in 1989, who was interested in its relationship to variations of the classical Oberwolfach Problem.

Question 4.3. For each positive integer ℓ, is there a positive integer d_ℓ such that every Eulerian graph with minimum degree d_ℓ admits an ℓ-locally self-avoiding Eulerian circuit? If so, what is d_ℓ for each ℓ?

The first part of the conjecture was verified for 4-edge-connected Eulerian graphs in 2016 [10], and then in general by Tien-Nam Le in 2017 [55]. That is, high minimum degree does turn out to be sufficient for guaranteeing the existence of locally self-avoiding Eulerian circuits. However, there is relatively little progress in pinpointing the values d_ℓ. The variant that we will investigate replaces Eulerian graphs of a specified minimum degree with other subclasses of Eulerian graphs, and then asks how high a girth we can guarantee. Specifically,

Question 4.4. Given a particular subclass $\mathcal{G}$ of Eulerian graphs, what is the greatest k such that all large enough graphs in $\mathcal{G}$ admit an Eulerian circuit that is $(k-1)$-locally self-avoiding?

Equivalently, this is asking for the smallest k such that every Euler circuit has a subcycle of length at most k. It is reasonable to allow for finitely many small counterexamples since we need enough room to detect the structural properties that force small subcycles. In the extreme case, for instance, we do not learn anything from the fact that a graph with k vertices cannot admit a k-locally self avoiding Eulerian circuit.

For a fixed Eulerian graph G, the value k in the previous question is the same as $g^E(G)$ in the notation of Oksimets [68], spread(G) used by Ramírez-Alfonsín [72], and the Eulerian recurrent length of G by Jimbo, Oshie and Hashiguchi [47]. The last set of authors showed that the problem of determining whether G admits an Eulerian circuit that is $(k-1)$-locally self-avoiding for a given k is NP-complete, and this is still true for the particular restricted case with G quartic and k greater than 330 [47, 48].

4.1.1 Known bounds

There are several classes of graphs that have been studied in relation to Question 4.4, with Eulerian graphs of maximum degree 4 being perhaps the earliest. The following forbidden subgraph characterisation of graphs in this class that admit a 3-locally self-avoiding (triangle-free) Eulerian circuit was proved by Adelgren [1] and independently by Heinrich, Liu and Yu [50].

Theorem 4.5. *A simple Eulerian graph with all degrees at most 4 admits a triangle-free Eulerian circuit if and only if it does not contain any of the subgraphs or configurations shown in Figure 4.1.*
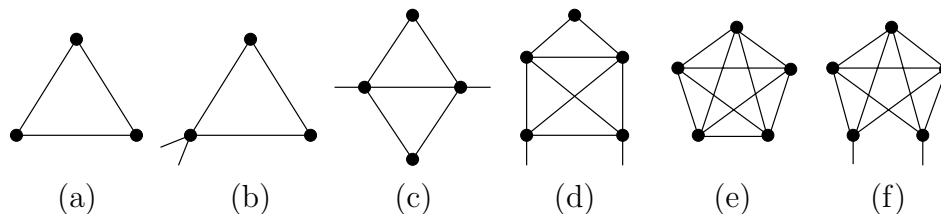

Figure 4.1: Obstructions to triangle-free Eulerian circuits.

Recall that as per our conventions for interpreting diagrams, the vertices shown do not have any more incident edges in the whole graph than are shown.

Both proofs proceed by induction on either the number of triangles, or the number of triangles together with the number of degree 4 vertices. As an outline, the idea is to remove a a 3-cycle or degree 4 vertex and then apply the induction hypothesis to obtain a triangle-free Eulerian circuit of the subgraph. The main work is then put to showing that this circuit can be rerouted to reincorporate those missing edges while avoiding all possible local conflicts; that is, avoiding all other possible triangles that might arise when we modify the circuit. This approach is, to some degree, the only thing we can do in a situation where the class of graphs is general enough that we do not know the exact structure of the graph that we are working with, so we leverage the degree bound by working locally.

More generally, Oksimets mentions a complete characterisation of simple eulerian graphs that admit triangle-free Eulerian circuits. In particular,

Theorem 4.6 (Oksimets [68])**.** *Every simple Eulerian graph G with minimum degree at least 6 admits a triangle-free Eulerian circuit. That is, for this class we have $k \geq 4$.*

Unfortunately, the proof of this result, as well as the statement of the full characterisation, appear only in theses that we have not been able to access.

The other two classes of graphs that have been studied are the complete graphs and complete bipartite graphs, which, in contrast to the classes we have seen so far, are very concrete. It therefore makes sense to let the subclass in the statement of Question 4.4 to be a single graph. In the latter case, we have the following exact answers:

Theorem 4.7 (Jimbo, [46])**.** *The greatest k such that $K_{m,n}$ (with $m \geq n$ both even) admits an Eulerian circuit that is $(k-1)$-locally self-avoiding is $2n-4$ if $n = m \geq 4$, and $2n$ otherwise.*

This agrees with an earlier solution in the special case for $K_{2n,2n}$ proved by Oksimets [68]. The case of complete graphs was first studied by Ramírez-Alfonsín [72]. The following theorem due to Oksimets gives the best current bounds.

Theorem 4.8 (Oksimets, [68])**.** *For the complete graphs K_{2n+1} with $n \geq 2$, $2n - 3 \leq k \leq 2n - 1$. As a special case, for K_7 we have $k = 4$.*

The same result were later proved independently by Jimbo [46], who also verified by computer that $k = 2n - 2$ for each $n \in \{3, 4, 5, 6\}$. We do not know of any theoretical proofs for the three larger values. Jimbo also conjectures that $k = 2n - 3$ for all larger complete graphs of odd order n.

The proof methods used in these are particular to complete and complete bipartite graphs as they rely on, for instance, the existence of edge-disjoint Hamilton circuits, or the fact that segments of consecutive vertices have a particular form (in the complete bipartite case) or are unrestricted (in the complete case). Additionally, these two classes are examples of cases for which the girth that can be guaranteed is quite high. Intuitively, this is due to the density and high connectivity which mean that if one was to try to construct a circuit without short subcycles, there are more options in terms of available edges and edge- or vertex-disjoint paths between vertices at each point, and hence more opportunities to avoid turning back on an already-visited vertex.

Our work will be on qp-graphs and tfp-graphs, which are closer in flavour to Eulerian graphs of maximum degree 4 in terms of structure and level of abstraction, and so it is on those arguments that we model our proofs. Nonetheless, between these classes we see the phenomenon discussed above whereby increasing connectedness also increases the attainable girth. In particular, we will show that if $\mathcal{G}$ is the class of quartic planar graphs, then the greatest k such that all large enough graphs in $\mathcal{G}$ admits a $(k-1)$-locally self-avoiding Eulerian circuit is 4. Restricting $\mathcal{G}$ to the 3-connected quartic planar graphs, we find that $5 \leq k \leq 8$.

4.1.2 Connection to edge-decompositions

Suppose we have a k-locally self avoiding Eulerian circuit of some graph G. Then provided $|E(G)|$ is divisible by some integer $\ell \leq k+1$, we can cut our circuit into

edge-disjoint segments each consisting of ℓ consecutive edges. The self-avoiding condition guarantees that the corresponding $\ell + 1$ vertices are all distinct, so each segment is isomorphic to $P_{\ell+1}$. This is called a $P_{\ell+1}$-decomposition of G. A similar construction given by cutting the circuit based on subcycles gives a cycle decomposition of G in which each cycle has length smaller than k.

These basic procedures illustrate the aforementioned relationship between k-locally self-avoiding Eulerian circuits and edge-decompositions of graphs, which broadly refers to a partition of the edges of a graph into components that satisfy some restrictions. A well-studied case is that in which all element of the partition induces a subgraph that is isomorphic to some fixed graph.

Definition 4.9. For two graphs G and H, we say that G has a H-decomposition if there is a partition $E_1 \cup \cdots \cup E_m$ of $E(G)$ such that the induced subgraphs $G[E_i]$ are isomorphic to H for all $1 \leq i \leq m$.

The trivial divisibility condition that $|E(H)|$ divides $|E(G)|$ is implicit in this definition. The following question posed by Bondy is concerned with the case that H is a path of fixed length.

Question 4.10 (Bondy [11])**.** Which simple graphs admit a P_k-decomposition?

This is similar to the path case of the Barát-Thomassen conjecture (see [3]). Every graph trivially has a P_1-decomposition, and a classical result due to Kotzig [54] says that a graph has a P_2 decomposition if and only if it has an even number of edges. For longer paths, results are limited to regular and complete bipartite graphs. A cubic graph has a P_3 decomposition if and only if it has a perfect matching [11, 54], and for r-regular graphs it is necessary and sufficient to satisfy the trivial divisibility condition and have no bridges when r is odd [50].

From Theorem 4.5, one can conclude by cutting a triangle-free Eulerian circuit into paths of length 3 that the Eulerian graphs of degree at most 4 all have P_3 decompositions. This is the key lemma in Heinrich, Liu, and Yu's proof of the following more general result.

Theorem 4.11 (Heinrich, Liu, Yu [50])**.** *A connected quartic graph admits a P_4 decomposition if and only if it has a 3-locally self-avoiding Eulerian circuit and $|E(G)|$ is divisible by 3.*

The other theorems from the previous section yield similar corollaries. We will give one more example.

Theorem 4.12 (Oksimets [68])**.** *The graph $K_{2n,2n}$ has a P_k-decomposition if and only if $k \le 4n - 1$ and $4n^2$ is divisible by k. Similarly, the graph K_{2n+1} has a P_k-decomposition if and only if $k \le 2n$ and $\binom{2n+1}{2}$ is divisible by k.*

We do not know of any other results for paths longer than length 4.

Note that the property of having a self-avoiding Eulerian circuit is stronger than having a P_{k+1}-decomposition. Given an appropriate path decomposition, there is no easy way to recover a locally self-avoiding Eulerian circuit since we cannot control subcircuits that occur across places where we have stitched paths together. Indeed, it may even be impossible. A corollary of our main theorem in this section is that all 3-connected quartic planar graphs have a P_5-decomposition. This includes the octahedron, which does not have a 4-locally self-avoiding Eulerian circuit.

4.2 A toy case

We now investigate Question 4.4 taking the quartic planar graphs as our particular subclass. Specifically, the statement is:

Question 4.4 for qp-graphs. What is the largest k such that all but finitely many quartic planar graphs admit an Eulerian circuit that is $(k-1)$-locally self-avoiding?

For Eulerian graphs with degree at most four, we have seen some positive results but not any complete answers. We have already mentioned the appeal of quartic graphs as the simplest nontrivial class of Eulerian graphs, and in light of Theorem 4.5, the quartic planar graphs seem to be a promising starting point since we have the following corollary.

Corollary 4.13. *Every quartic planar graph admits a 3-locally self-avoiding Eulerian circuit.*

Proof. There is nothing more to do than verify that the quartic planar graphs cannot have any of the forbidden subgraphs shown in Figure 4.1. Of those, only (e) and (f) can be quartic. We can rule out (e) immediately as it is isomorphic to K_5. For (f), suppose we have a quartic planar graph G that contains $K_5 - xy$ as a subgraph. Denote the rest of the graph by H, and let a and b be the unique neighbours of x and y respectively that are in H. Note that H is connected, for otherwise $\{a\}$ would be a 1-cut in G, and the components of $G - a$ would

each have an odd number of odd degree vertices. This means that there is an ab-path in H, and consequently G contains a subdivision of K_5 which contradicts planarity. □

Returning to the main question, this shows that $k \geq 4$ since we can always avoid triangles. In fact, this is sharp.

Theorem 4.14. *Let G be any quartic graph containing the diamond gadget shown in Figure 4.2. Then every Eulerian circuit in C has girth at most 4.*
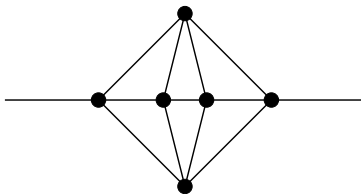

Figure 4.2: The diamond gadget forces subcycles of length at most 4.

Since the gadget has two vertices of attachment (which need not be distinct), it can only be entered once. That is, if an Eulerian circuit through some graph containing this subgraph is to have girth 5 or greater, then there must be a path through the diamond gadget that avoids 3- and 4-cycles. It is tedious but straightforward to verify that no such path exists. One can easily construct an infinite family of quartic planar graphs containing this gadget, so we therefore have the following result.

Theorem 4.15. *The answer to Question 4.4 for quartic planar graphs is 4.*

4.3   Imposing 3-vertex-connectedness

Following the last result, the natural extension is to consider 3-connected quartic planar graphs since 3-connectedness immediately excludes all graphs in the family of counterexamples from the previous section. After this refinement, our object is now to answer Question 4.4 for tfp-graphs:

Question 4.4 for tfp-graphs. What is the largest k such that all but finitely many 3-connected quartic planar graphs admit an Eulerian circuit that is $(k-1)$-locally self-avoiding?

Although the diamond gadget is no longer admissible since the vertices of attachment would form a 2-cut, Faller and McKay [32] have found a family of

graphs formed from a gadget with four vertices of attachment that gives us an upper bound on k.

Theorem 4.16. *Let $\mathcal{G}$ be the family of 3-connected quartic planar graphs formed by joining an odd number of copies of the gadget shown in Figure 4.3 end to end in a cycle. Then for all $G \in \mathcal{G}$, every Eulerian circuit in G has girth at most 8.*
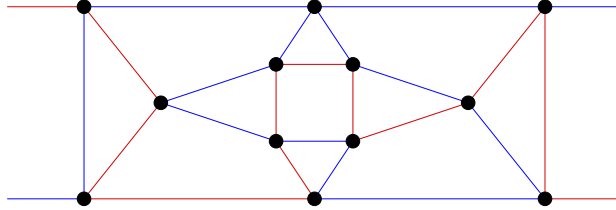

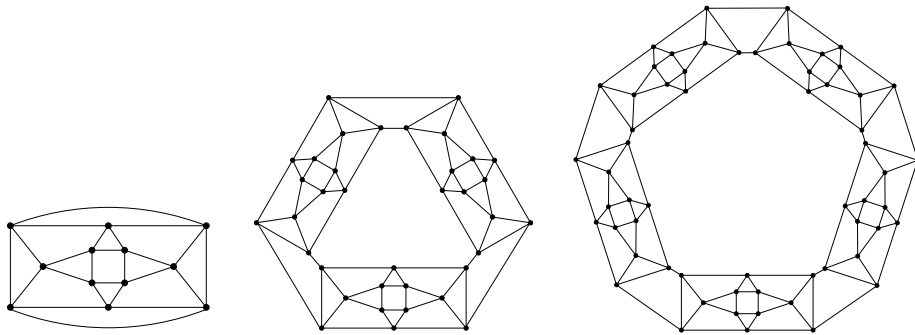Figure 4.3: A gadget that forces subcycles of length at most 8.

Figure 4.4: Three members of $\mathcal{G}$.

The three smallest members of $\mathcal{G}$ are shown in Figure 4.4. Since there are four edges going in or out of the gadget, it is convenient to view the Euler circuit as making two passes through, which we represent by the different colours. The last theorem was verified by testing all possible ways to traverse the edges of the gadget by computer, which showed that in order to avoid subcycles of length at most 8, each pass through the gadget must enter at the top left and exit a lower right vertex, or else enter at the lower left and exit at the top right where the relative positions are with reference to the drawing in Figure 4.3. Joining an odd number of such paths together forces at least one copy to be traversed with paths that enter and exit through the two upper or lower vertices, and hence must have a subcycle of length at most 8.

On the other hand, Corollary 4.13 still applies to the 3-connected case so we have an immediate lower bound of 4. The remainder of this section is devoted to proving the following theorem which improves the lower bound to 5.

Theorem 4.17. *Every 3-connected quartic planar graph except the octahedron has a 4-locally self-avoiding Eulerian circuit.*

Once this is proved, we will have shown that $5 \leq k \leq 8$ for 3-connected quartic planar graphs. Theorem 4.17 is also independently interesting, as there have so far not been any results on avoiding 4-cycles. As the proof is somewhat involved, we will first discuss the strategy and set up some theory before diving in. For brevity, following [23] we will use the term *tfp-graph* to refer to 3-connected quartic (4-regular) planar graphs.

The plan of attack is to follow the general outline of the proofs of Theorem 4.5. In particular, we would like to use induction since our class of graphs is broad and does not come with a description that would allow us to construct locally self-avoiding Eulerian circuits for each member from scratch. Note that we do not have the luxury of being able to removing a small cycle or a vertex to apply the induction hypothesis, since in general this will not preserve 3-connectedness, nor will the resulting graph be quartic. This means we need some cleverer way to reduce a given 3-connected quartic planar graph while preserving all of the properties defining this class. The generation theorems discussed in the previous chapter provide just that.

For this application, it is advantageous to not only have expansion operations that are as restrictive as possible, but also to have a range of operations that will enable us to apply reductions in more situations. Specifically, we will use the following with elements borrowed from generation theorems of [17] and [23].

Corollary 4.18. *The class of 3-connected quartic plane graphs can be generated from the antiprisms by (restricted) pegging, 4-cycle addition, and 3-cycle slides depicted in Figure 4.5.*

It is immediate that this is a generation theorem since all we have done is taken the dual of Theorem 3.6(3) and added in one superfluous operation. There is no harm in this, since the third operation comes from the generation theorem in [23] and hence we can be sure that its inverse preserves the required properties of the class. In all, our three operations are the duals of E_1, E_3 and a less restricted version of E_2.

Remark 4.19. The pegging operation shown above is specifically a *restricted* pegging, meaning it is only applied when the edge *bd* is present. This reduces the number of cases we need to consider later on. It is also interesting to note that there is a precedent for pegging in the study of Eulerian circuits, set by Tutte and Smith in 1941 [85].
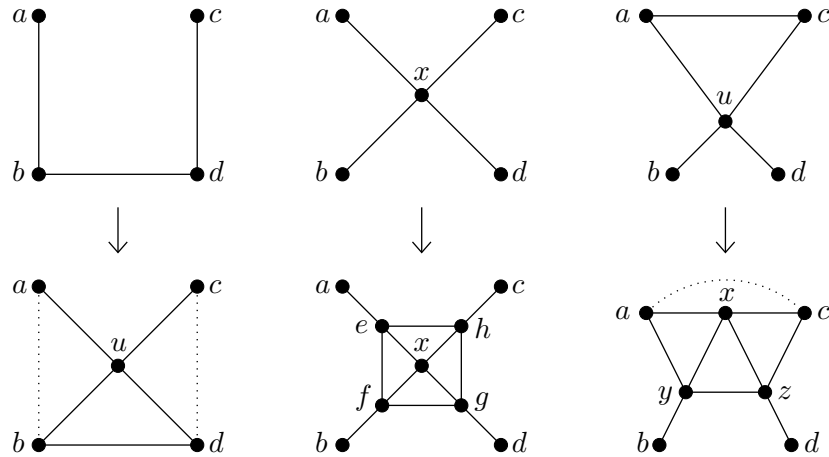
Figure 4.5: From left to right, these depict pegging, 4-cycle addition and 3-cycle slide. Dotted lines indicate edges in the complement.

To run the induction argument, we will need to show that each of the starting graphs have a 4-locally self-avoiding Eulerian circuit, and that each of these operations preserves the property of there existing such a circuit. The proof will span the next three sections. We first handle 4-cycle addition and 3-cycle slides (Section 4.3.1), and then tackle the considerably more involved case of pegging (Section 4.3.2), before piecing the argument together (Section 4.3.3). Throughout, we will freely use the fact that 3-connected graphs have a unique embedding on the sphere, and assume a fixed choice of outer face. In addition, we will refer to 4-locally self-avoiding Eulerian circuits as simply *good* circuits, and define a *short* cycle to be one of length 3 or 4.

4.3.1   Circuit extensions

For the operations of 4-cycle addition and 3-cycle slides, it turns out to be easy to extend an induced Eulerian of the smaller graph to cover the new edges and vertices in the expanded graph.

Lemma 4.20. *Suppose G and H are tfp-graphs such that G can be constructed from H by a single 4-cycle addition. If H admits a 4-locally self-avoiding Eulerian circuit, then G does as well.*

Proof. We will show that the 4-locally self-avoiding Eulerian circuit in H can be extended to include the new edges created in the expansion whilst still avoiding 3- and 4-cycles. After the edge ax appears in the circuit, we must traverse one of bx, cx or dx. Suppose we have axb, so that the blue edges in the rightmost diagram

in Figure 4.6 are consecutive. This forces the red edges to also be consecutive in the circuit, so cxd (or equivalently, dxc). To extend the circuit, we replace axb by the path $aexhgfb$, and cxd by $cgxfehd$ as indicated in the figure. The resulting circuit is certainly Eulerian, and we have not created any short subcycles within either of the paths. There is also no chance of creating short cycles nearby; if we identify ax, bx, cx and dx in H with ae, bf, cg and dh in G respectively, then it can be observed that any subcycle involving x increases in length when we extend the circuit. The case that we have axd is symmetric, and an appropriate extension can be obtained by rotating the figure.

If instead we have axc in the original circuit, then we must also have bxd. We replace these paths by $aefxhgc$ and $bfgxehd$ respectively. Again, the resulting circuit is Eulerian, and by the same argument as before it does not have any short subcycles. $\qquad\square$
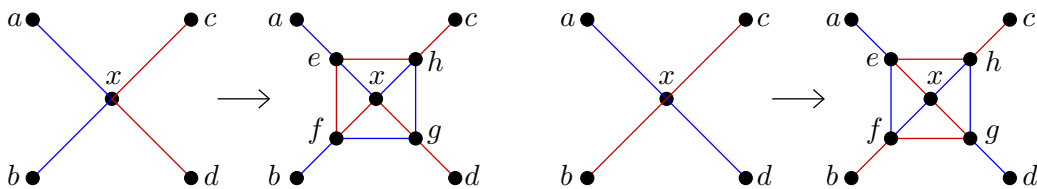
Figure 4.6: Extension of an Eulerian circuit after a 4-cycle addition.

Lemma 4.21. *Suppose G and H are tfp-graphs such that G can be constructed from H by a single 3-cycle slide. If H admits a 4-locally self-avoiding Eulerian circuit, then G does as well.*

Proof. Suppose we have a 4-locally self-avoiding Eulerian circuit in H. At vertex u, the circuit can either proceed straight ahead or turn. The two rows in Figure 4.7 represent these two situations, where edges of the same colour are assumed to appear consecutively. We allow that the blue edge ad may or may not be consecutive to either the red or green edges.

If auc is in the original circuit, to obtain an Eulerian circuit in G we make the following substitutions; auc becomes $ayzc$, bud becomes $byxzd$, and ad becomes axd. These paths are shown in the middle picture of the top row of Figure 4.7. The only way a short subcycle may occur is if ad is consecutive to the green edges, causing the edges of the triangle xzd to be traversed consecutively. In this case, we use the alternative substitutions shown in the rightmost picture of the top row; auc becomes $ayxzc$, bud becomes $byzd$, and ad becomes axd. We would

have another problematic triangle *axy* if *ad* were to occur directly next to the red edges in the circuit, but this is impossible since if *ad* were adjacent to both the red and green edges, we must have *uadu* as a subcycle of the original circuit which was assumed to be 4-locally self-avoiding.

Similarly, if *aub* is in the original circuit, we obtain an Eulerian circuit in *G* by one of the following substitutions; either *aub* becomes *ayb*, *cud* becomes *czyxd*, and *ad* becomes *axzd*, or we replace *aub* by *axzyb*, *cud* by *czd*, and *ad* by *ayxd*. In either case, no 4-cycles can be created. The former yields a circuit that will only have a 3-cycle if *ad* is consecutive to the green edge, whilst the latter can only produce a 3-cycle if *ad* is consecutive to the red edges. We have already shown that these cannot both be true simultaneously, so one of the resulting circuits is free of short cycles. □
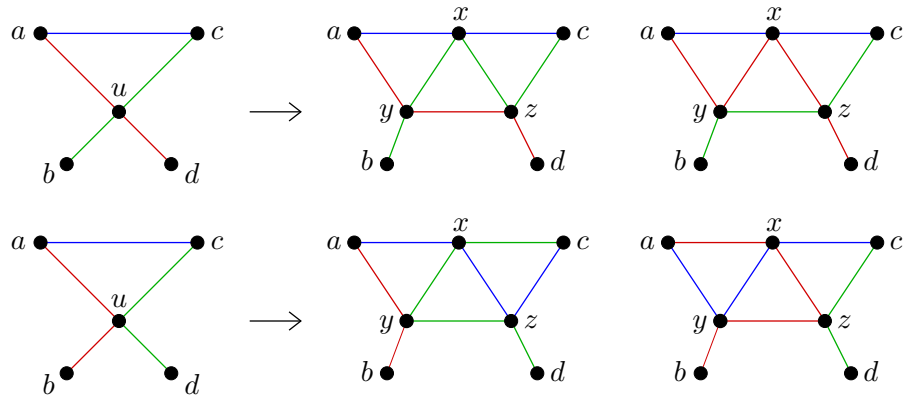
Figure 4.7: Extension of an Eulerian circuit after a 3-cycle slide.

4.3.2 Rerouting

The two expansion operations we have handled so far have admitted simple circuit extensions largely because they involve creating multiple new vertices, and we know precisely what the neighbours of each of these are. Moreover, all of the vertices in the smaller graph that are affected by the expansion were already connected by short paths, so in most cases the expansion merely lengthens each cycle. Unfortunately, the pegging operation only produces one new vertex, so there is very little room for extension. Instead, the technique we use is to start with an induced circuit which may have short subcycles, and then resolve these by rerouting the circuit locally. It must be verified that the modified circuit does not create any new problematic subcycles, or if it does, that those can also be resolved by rerouting again.

To perform this rerouting, we generally need to split the circuit obtained by the induction hypothesis into segments, which we denote by uppercase letters W, X, Y, Z. A segment X from vertex a to vertex b means that aXb is a subpath in our circuit, where $X = x_1 x_2 \ldots x_k$ is said to contain k vertices and have length $k + 1$, so we consider it to include the edges ax_1 and $x_k b$. Although we use vertex notation for convenience, we are of course ends concerned with edges. For instance, the inverse segment $X_{-1} := x_k x_{k-1} \ldots x_1$ starts with the edge bx_k and ending with the edge $x_1 a$. For vertices at the end of the segment, we will also use the notation x_{-1}, x_{-2} for the last and second last vertices respectively. It is possible for a segment to contain no vertices, in which we say it is just an edge.

The following example illustrates our notation and the technical part of our method. We will not be quite so explicit again after this.

Example 4.22. Suppose we start with the 5-antiprism labelled as shown in Figure 4.8. A 4-locally self-avoiding circuit is given by $abcdeafgdhcibjfeghija$. If we peg the edges ij and hg to create a new vertex k, then there is an induced Eulerian circuit obtained by replacing ij in our previous circuit with ikj and gh with gkh. The induced circuit is therefore $abcdeafgdhcibjfegkhikja$. Notice that this now has a subcircuit $khik$ of length 3, bounding the shaded region.

Let us rewrite the induced circuit as $khikXhYiZk$ where the segments are $X = jabcdeafgd$ shown in red, $Y = c$ in green and $Z = bjfeg$ in blue. To resolve the triangle, we reroute to $khiZkiY^{-1}hX^{-1}k$ which we see is indeed an Eulerian circuit. The reader may like to verify that it is 4-locally self-avoiding. To be helpful, the final circuit is $ajkhibjfegkichdgfaedcba$.
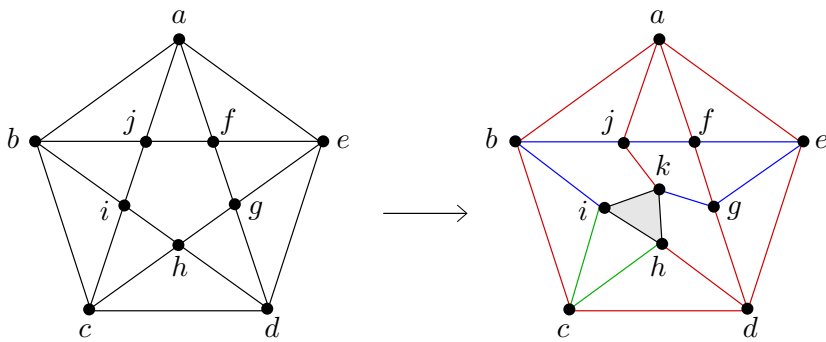

Figure 4.8: Pegging and rerouting.

Lemma 4.23. *Let G be a 3-connected quartic planar graph not isomorphic to the octahedron, and suppose it has an Eulerian circuit C. Then each vertex in G can lie on at most one 3-cycle or 4-cycle in C.*

Proof. Since G is quartic, each vertex occurs twice in the circuit except the initial vertex which appears three times. It is also clear that each vertex can be in at most 2 cycles. Suppose $a \in V(G)$ is part of two 3-cycles *abca* and *abda*, and write the circuit so that it is the initial and terminal vertex. In order for a to appear three times, these 3-cycles must be traversed back-to-back so that we have *abcabda*. Moreover, since the circuit is connected by definition, this must be all of C so we see that G has 5 vertices. In the case that a is on one 3-cycle and one 4-cycle, or two 4-cycles, the same argument shows that G can have at most 6 or 7 vertices respectively. In all cases, we have a contradiction since the smallest 3-connected quartic planar graph not isomorphic to the octahedron has 8 vertices. $\qquad\square$

Lemma 4.24. *Suppose we have a connected plane graph W with two vertices of degree 1, one vertex of degree 2 and all other vertices have degree 4. Suppose further that the vertices of degree less than 4 are pairwise non-adjacent and there is some face incident to all of them. Then W has at least 5 vertices of degree 4.*

Proof. Let the degree 1 vertices be p and q, and the degree 2 vertex be r. If none of p, q and r have neighbours in common then by identifying them we obtain a simple quartic planar graph. The smallest possibility is the octahedron which has 6 vertices. Hence, in W there must be at least 5 vertices of degree 4.

Otherwise, $|N(p) \cup N(q) \cup N(r)|$ is equal to 2 or 3. If it is 3, then let the set of their neighbours be $\{r, s, t\}$. Two of these are adjacent to exactly one vertex in $\{p, q, r\}$. Focussing on one of them, as it has degree 4, it must have another neighbour u distinct from all of $\{p, q, r, s, t\}$. Similarly, of the named vertices u can be adjacent to at most r, s and t meaning it, too, has another neighbour v not in the set $\{p, q, r, s, t, u\}$. Now r, s, t, u, and v are the required vertices.

Suppose instead that $N(p) \cup N(q) \cup N(r) = \{r, s\}$. In this case, among the neighbours of r and s there must be at least one new vertex t of degree 4. Since W is simple and the neighbours of p, q and r are already known, then $N(t) = \{r, s, u, v\}$ where u and v are both vertices of degree 4 and not in $\{p, q, r, s\}$. That already makes five vertices of degree 4. $\qquad\square$

For our purposes, it would have been enough to require three vertices of degree four in the previous lemma, but five is a sharp lower bound.

There is one other notion that we will use, and that is of a *trapped* vertex. The idea is that in the course of resolving local conflicts, we often find short cycles where the cyclic ordering at some of the vertices is known, and which perhaps is

adjacent to a facial cycle. This can sometimes produce a degree-deficient vertex in our quartic graph that cannot be included in the Eulerian circuit without either creating a 2-cut or violating planarity.

Example 4.25. Figure 4.9 depicts two representative traps. Suppose in our construction we have the black subgraph on the left of the figure with cyclic orderings as shown, and all other edges of the whole graph are in the a single segment E that starts at u and ends at b. Clearly, be_1 is outside the region bounded by $uvwcu$. In order for E to contain vertex w, the edge $x_{-1}u$ must be inside that cycle since there are no other degree-deficient vertices that act as entry points into this bounded region. However, as there is no exit point, this means that E cannot be connected which is a contradiction. This is the sense in which w is trapped.

In the second case, we have introduced a vertex w which means that E can actually enter $uvwcu$. To avoid the previous contradiction, we are now forced to have the cyclic orderings of black edges and E-edges as drawn and we see that $\{b, x\}$ is a 2-cut. In this situation, we still say that w is trapped but we will also mention that a 2-cut is created.
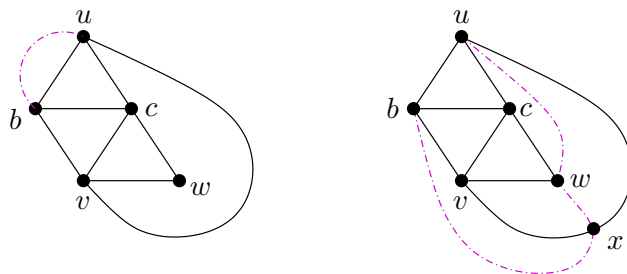

Figure 4.9: Trapped vertices.

We offer two words of warning in relation to interpreting diagrams in this section. First, we do not care about the order of half-edges that are in segments of Eulerian circuits if they occur consecutively in the cyclic ordering at a vertex. This means that at w in Figure 4.9, the edge occurring directly after cw may be closer to u or to x in E. Secondly, the pair $\{u, w\}$ is not necessarily a 2-cut since that segment of E cannot be assumed to be vertex-disjoint with the segment between w and x. We draw it as such to avoid clutter, and will continue to do so in the diagrams to come.

Lemma 4.26. *Suppose G and H are tfp-graphs such that G can be constructed from H by a single (restricted) pegging. If H admits a 4-locally self-avoiding Eulerian circuit, then G does as well.*

Proof. Let the expanded graph G be obtained from H by pegging the edges ab and cd (refer to picture 1). Since we only require restricted pegging, we may assume that the edge bc is also present. Since H is also a 3-connected quartic planar graph (different from the octahedron), by the induction hypothesis there is an Eulerian circuit $XabYcdZ$ in H that is 3- and 4-cycle-free. After pegging, we have an induced Eulerian circuit in G given by $XaubYcudZ$.

Let's examine the situations in which this induced circuit is not 4-locally self-avoiding. Any subcycle in the circuit that does not contain the vertex u was also present in the original circuit in H, and hence cannot be a 3-cycle or 4-cycle. Therefore, any short cycle must contain u. On the other hand, lemma 4.23 says that u can lie on at most one 3-cycle or one 4-cycle, so together these imply that the induced circuit contains at most one short cycle. We can say explicitly what these problem cycles must look like.

- If u is part of a 3-cycle, then since the edges ab and cd are not in G by the definition of pegging, the only possibilities are $ubcu$ and $uadu$. These are symmetric, so we will only handle the former.

- If u is part of a 4-cycle, then there are four possible forms; $ubvcu$, $uavdu$, $ubvdu$ and $uavcu$ where v is some vertex in G distinct from a, b, c, d and u. Here, the first two are symmetric and the last two are symmetric. Note that $uavbu$ and $ucvdu$ are not valid, since in H they would correspond to the cycles $avba$ and $cvdc$.

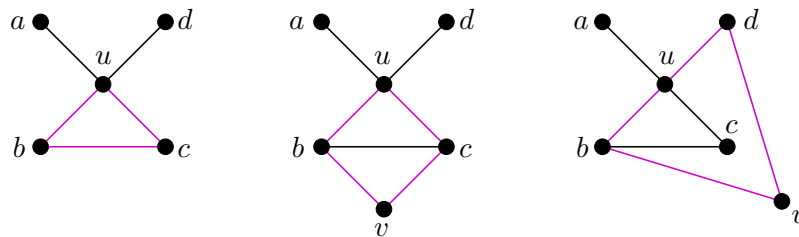For each of these three types of short cycle, we will show that they can be avoided by rerouting the circuit.

Figure 4.10: Three types of problem cycles, shown in magenta.

Case 1: Resolving a 3-cycle of the form ubcu

We begin by finding a convenient general form of the induced circuit. To do this, we break it into three segments X (between u and b), Y (between b and c) and Z (between c and u). The circuit can then have the form $ubcuXbYxZu$ or $ucbuXbYxZu$. Assume that we have the former - we will eventually see that the following argument goes through just as well for the latter. Refer to picture 2. We split into three subcases depending on whether one, two or all of the segments have at least 3 vertices. Note that it is not possible for all three segments to be so short; the smallest tfp-graph that is not the octahedron has 8 vertices, so the average number of vertices on each segment is $\frac{10}{3} > 3$.

Case 1a. In the case that X, Y and Z all have at least 3 vertices, we can try any of the symmetric circuits shown in Figure 4.11. There are actually six circuits of this form, however we exclude those in which X and Z appear consecutively since these segments are separated by the triangle ubc in the original circuit, leaving us with these four possibilities.

$$ucbX^{-1}ubYcZu \qquad ucbYcZubX^{-1}u \qquad ubcZucY^{-1}bX^{-1}u \qquad ubcY^{-1}bX^{-1}ucZu$$
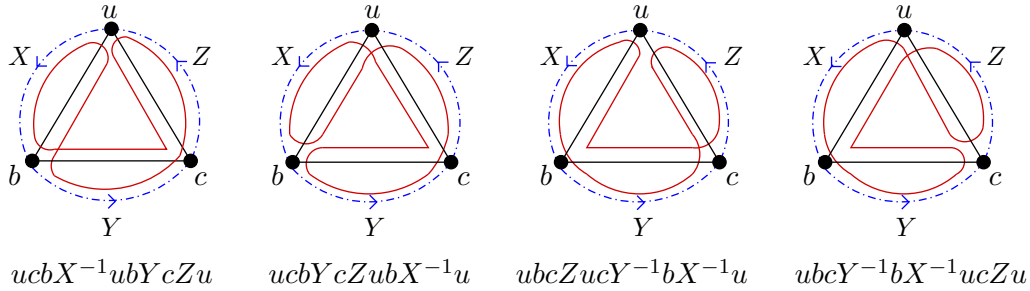
Figure 4.11: Possible Eulerian circuits when each segment is long.

The possible short subcycles that may arise for each of these circuits are enumerated below.

1. For $ucbX^{-1}ubYcZu$, we have a 3-cycle if $x_1 = y_1$ and a 4-cycle if $x_1 = y_2$ or $x_2 = y_1$.

2. For $ucbYcZubX^{-1}u$, we have a 3-cycle if $x_{-1} = z_{-1}$ and a 4-cycle if $x_{-1} = z_{-2}$ or $x_{-2} = z_{-1}$.

3. For $ubcZucY^{-1}bX^{-1}u$, we have a 3-cycle if $y_{-1} = z_{-1}$ and a 4-cycle if $y_{-1} = z_{-2}$ or $y_{-2} = z_{-1}$.

4. For $ubcY^{-1}bX^{-1}ucZu$, we have a 3-cycle if $x_1 = z_1$ and a 4-cycle if $x_1 = z_2$ or $x_2 = z_1$.

It turns out that none of the 3-cycles mentioned above can really occur under our assumptions. Recall that vertices and a and d are also adjacent to u, and in particular $\{a, d\} = \{x_1, z_{-1}\}$. Suppose $a = x_1 = y_1$, then we would have the edge $y_1 b = ab$ in G, but this contradicts the definition of pegging. If instead $d = x_1 = y_1$, then due to fixed cyclic ordering of vertices at u, we have one of the situations shown in Figure 4.12, drawn so that vertex a is inside the triangle ubd. From left to right, the first situation has a 2-cut $\{a, b\}$, the second has a 2-cut $\{c, d\}$ and the third has a 1-cut $\{c\}$ so in all cases we have a contradiction to G being 3-connected. Thus, we conclude that $x_1 \neq y_1$ so there are no 3-subcycles in $ucbX^{-1}ubYcZu$. An identical argument also rules out the possibility that $x_{-1} = z_{-1}$, $x_{-1} = y_{-1}$ or $x_1 = z_1$.
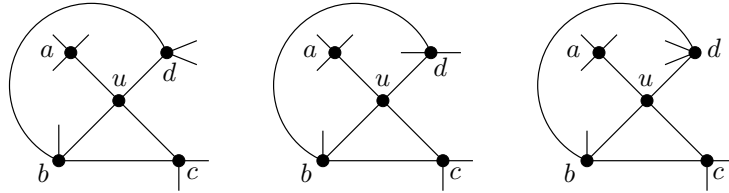
Figure 4.12: Possible cyclic orderings at a, b, c and d if $d = x_1 = y_1$.

Since G is quartic, each vertex not in $\{a, b, c, d, u\}$ appears on exactly two of X, Y and Z. This means that we cannot, for instance, have both $x_1 = y_2$ and $x_1 = z_2$. Suppose we try both $ucbX^{-1}ubYcZu$ and $ubcY^{-1}bX^{-1}ucZu$. In order for both of these to have short subcycles, it must be that $x_1 = y_2$ and $x_2 = z_1$, or $x_1 = z_2$ and $x_2 = y_1$. We treat these separately.

1ai. Suppose that $x_1 = y_2$ and $x_2 = z_1$. The only possible cyclic ordering of edges that does not imply the existence of a 2-cut is that shown in the leftmost picture of Figure 4.13; if uby_1x_1u is nonfacial, then Y_1 must lie on Z but this means that $\{u, y_1\}$ is a 2-cut, and if uby_1x_1u is facial then either Y is contained inside $bcz_1x_1y_1b$ in which case $\{u, z_1\}$ is a 2-cut, or we have the configuration shown. This embedding implies that X and Z are disjoint, since they are contained in opposite regions bounded by the cycle $by_1x_1x_2cb$. It follows that $ucbYcZubX^{-1}u$ is a good Eulerian circuit as neither $x_{-1} = z_{-2}$ nor $x_{-2} = z_{-1}$ can hold.

1aii. Suppose that $x_1 = z_2$ and $x_2 = y_1$. If ux_1x_2bu is facial, then z_2z_3 must occur between uz_2 and z_1z_2 in the cyclic ordering at u, otherwise $\{u, z_1\}$ would be a 2-cut. Then if z_1 lies on Y as well as Z, or appears on Z twice,

this forces the remaining vertices to have cyclic orderings shown in middle
two configurations of Figure 4.13. In both cases, X and Z are disjoint, so
$ucbYcZubX^{-1}u$ is a good Eulerian circuit. Note that z_1 cannot lie on X,
since then $\{u, x_1\}$ would be a 2-cut. If ux_1x_2bu is non-facial, then to avoid
2-cuts we are forced to have the cyclic ordering shown in rightmost picture
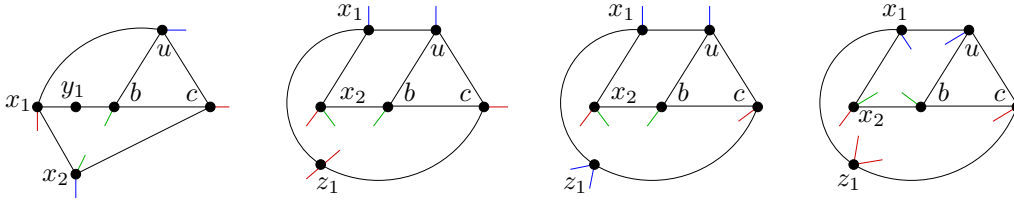of Figure 4.13. Now Y and Z are disjoint so the circuit $ubcZucY^{-1}bX^{-1}u$
will do.

Figure 4.13: Cyclic orderings for cases 1ai and 1aii. Green half-edges are on X,
red on Y and blue on Z.

Case 1b. Suppose two of X, Y and Z have at least 3 vertices, and the third has
at most 2 vertices. Let us first assume that X and Z are the longer segments.
We again consider the circuit $ucbX^{-1}ubYcZu$, shown in the leftmost diagram
of Figure 4.11 from the previous case, which we recall will be a good Eulerian
circuit unless $x_1 = y_2$ or $x_2 = y_1$ and that $x_1 \neq y_1$ is guaranteed by the argument
from (a). Combining this with the possible number of vertices in Y leads to three
subcases.

1bi. If $x_1 = y_2$, then $|Y| = 2$. Write $X = x_1X'$. Suppose az_{-1} is in the region
bounded by ux_1cbu that does not contain y_1. Then all of X must lie in the
opposite region as well, in which case $\{u, c\}$ is a 2-cut. Now suppose instead
that az_{-1} is in the region bounded by ux_1cbu that also contains y_1. Since
the initial vertex of Z is adjacent to c, it must be that y_2 is on both Z and
Y. In particular, it is not on X, which must therefore be contained in either
ux_1y_1bu or by_1x_1cb. Then either $\{u, y_1\}$ or $\{c, y_1\}$ is a 2-cut. Examples of
these configurations are represented in the first column of Figure 4.14.

1bii. If $x_2 = y_1$ and $|Y| = 1$, then it is clear that by_1c must be facial in order
to avoid 2-cuts. This determines the cyclic ordering of edges at c. Suppose
uz_{-1} is between ux_1 and uc in the ordering at u. If x_1 lies on Z, which is
the situation shown in the first picture of the second column of Figure 4.14,
then $\{x_1, b\}$ is a 2-cut. Similarly, if the vertex x_1 appears twice on X,

then $\{x_2, b\}$ is a 2-cut. Otherwise, uz_{-1} occurs between ux_1 and ub in the ordering at u. Then x_1 must lie on Z, and $\{x_1, c\}$ is a 2-cut which may be observed from the second picture.

1biii. If $x_2 = y_1$ and $|Y| = 2$, we first observe that we cannot have uz_{-1} be between ux_1 and ub in the cyclic ordering at u and cz_1 between bc and y_2c in the cyclic ordering at c simulateously. Since this situation is symmetric, we may assume without loss of generality that cz_1 is between uc and y_2c. Letting $X = x_1x_2X'$, if X' has vertices in both regions bounded by the cycle ux_1x_2bu then $\{x_1, b\}$ is a 2-cut. Otherwise, X' in the region bounded by $\{x_1, b\}$. The right column of Figure 4.14 depicts the two possible situations, and, for a change, neither of these lead to a contradiction. Rather, for the upper configuration, we may take the circuit $ubx_2y_2cux_1x_2X'bcZ'x_1Z''u$ where we have split $Z = Z'x_1Z''$. Here, Lemma 4.24 ensures that Z' is long enough that there cannot be a short subcycle in $X'bcZ'x_1Z''$. For the lower configuration, the circuit $ucy_2x_2x_1uZ^{-1}cbx_2X'bu$ is 4-locally self-avoiding, again using Lemma 4.24 to see that $bx_2X'b$ is not a short subcycle.
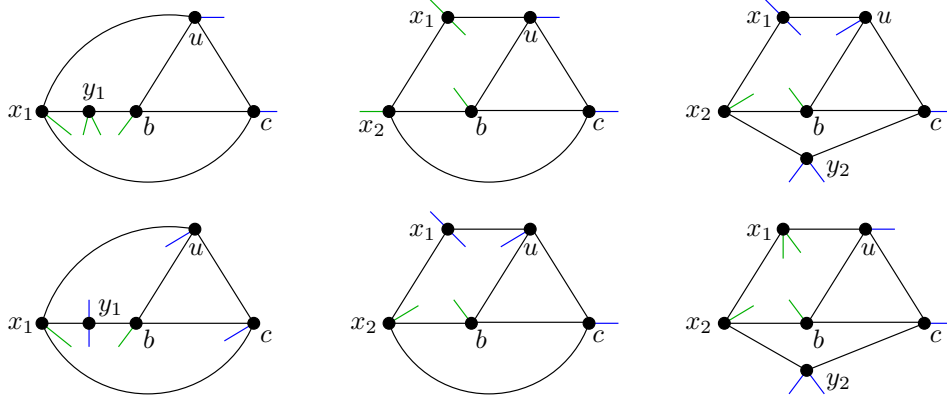

Figure 4.14: Cases 1bi, ii and iii. Green half-edges are on X', and blue on Z.

If X and Y are the longer segments, we can use the circuit $ucbY cZubX^{-1}u$, and for Y and Z we use $ubcY^{-1}bX^{-1}ucZu$, both of which we have already seen in 1a. These are known to be 3-cycle free already, and an appropriate rotation of the above argument ensures that they are also 4-cycle free.

Case 1c. Since ab and cd are not edges in G, and bc is in the cycle, it is immediate that X and Z have at least 2 vertices, and Y has at least 1 vertex. Without loss of generality, suppose that X and Y are the short cycles.

Suppose that $X = x_1 x_2$ and $Y = y$. The vertices x_1, x_2 and y are necessarily distinct; certainly $y \neq x_2$ since they are both adjacent to b, and setting $y = x_1$ would trap x_2. Thus, we have the general configuration shown in Figure 4.15. Observe that $ubcu$ and $cbyc$ are necessarily facial, else there would be a 2-cut. If $ux_1 x_2 bu$ is non-facial, then we can further divide Z into two segments; let Z' be the segment from c to either x_1 or x_2, and Z'' the segment from either x_1 or x_2 to u. Here, we are splitting at the unique vertex out of x_1 and x_2 that is incident to one edge outside of $ux_1 x_2 bu$ and one inside the same cycle. Assume we split at x_2 (there are analogous circuits if we split at x_1). Then $uycZ'x_2 x_1 ucbx_2 Z''u$ is a good Eulerian circuit unless $Z_2' = y$ or $Z_3' = y$. If $Z_2' = y$, then we have the three-triangle configuration. Otherwise, let $Z' = z_1 z_2 Z'''$. Then $ubcyZ'''x_2 x_1 ucz_1 z_2 ybx_2 Z''u$ is a good Eulerian circuit.
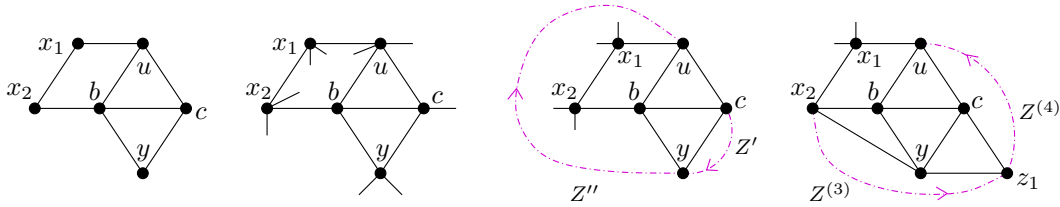

Figure 4.15: Configurations in case 1c.

If $ux_1 x_2 bu$ is facial, let us split Z at y and write $Z = Z'yZ''$ depicted in the third diagram of Figure 4.15. Consider the circuit $ux_1 x_2 bcyZ''ubyZ'^{-1}cu$. There are a few potential short subcycles to check here. Firstly, if $Z_1'' = x_2$, then since ux_2 is not an edge (that would trap x_1, given the assumption that $ux_1 x_2 bu$ is facial), either we can apply a 3-cycle unslide centred at b in which case we are done, or Z'' has only one vertex, say z_1. If the latter is true, we can split again at z_1 to get $Z' = Z^{(3)} z_1 Z^{(4)}$. Then $ux_1 x_2 bcz_1 Z^{(3)} x_2 ybuZ^{(4)} z_1 ycu$ is a good Eulerian circuit unless $Z^{(4)}$ is a single edge, in which case $ux_1 x_2 yz_1 ucbx_2 Z^{(3)-1} z_1 cybu$ suffices. Otherwise, if $Z_1'' \neq x_2$ then there is no short subcycle in $x_2 bcyZ''$.

The remaining checks are much easier. For sake of brevity, we simply note that the potential conflicts occur if Z'' is an edge, $Z_1' = x_1$, $Z_1' = x_2$, Z'' has exactly one vertex, $Z_1' = x_1$ or $Z_{-1}'' = Z_{-1}'$. The first of these traps x_1 and x_2, whilst the others create 2-cuts so we reach contradictions in all cases.

The case that $X = x$ and $Y = y_1 y_2$ can be handled using symmetric circuits, so the remaining case is that $X = x_1 x_2$ and $Y = y_1 y_2$. This is depicted in the leftmost picture of Figure 4.16. These four internal vertices must all be distinct to avoid trapping vertices; it is straightforward to verify that if $x_1 = y_1$, $x_2 = y_2$ or $x_1 = y_2$, each of which is shown in Figure 4.16. As before, $ubcu$ is necessarily

facial. Suppose that both ux_1x_2bu and by_1y_2cb are non-facial. Then Z splits as $Z = Z'y_iZ''x_jZ'''$ where $i, j = 1$ or 2, and Z' and Z''' are contained respectively in the faces ux_1x_2bu and by_1y_2cb. However, this means that $\{x_j, y_i\}$ is a 2-cut, so this configuration cannot occur.
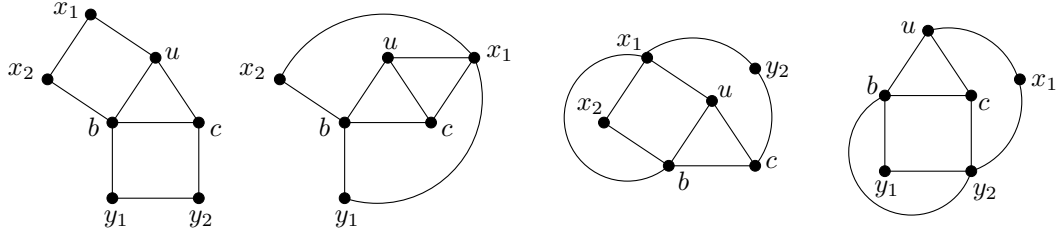

Figure 4.16: Possibilities when x_1, x_2, y_1 and y_2 are not all distinct.

If ux_1x_2bu is facial but by_1y_2cb is non-facial, we may write $Z = Z'y_iZ''$ with $i = 1$ or 2 so that Z' is contained inside by_1y_2cb. Then $uZ''^{-1}y_1y_2cux_1x_2by_1Z'cbu$ is a good Eulerian circuit unless Z'' is an edge. This cannot occur since x_1 and x_2 would subsequently be trapped. We have also used that Z' has at least 5 vertices, which follows from Lemma 4.24.

When ux_1x_2bu and by_1y_2cb are both facial, we reroute to $ux_1x_2bcZuby_1y_2cu$. Note that Z must contain x_1, x_2, y_1 and y_2 so $bcZub$ will not be a short cycle. The only way a short subcycle can be present here is if $z_1 = x_2$, $z_2 = x_2$, $z_{-1} = y_1$ or $z_{-2} = y_1$. The first and third of these trap either x_1 or y_2. If the second case holds, the situation is as shown in rightmost diagram along Figure 4.17. Here, a good circuit is given by $uby_1y_2cz_1Z''ucbx_2Z'z_1x_2x_1u$, where we have split Z as $z_1x_2Z'z_1Z''$. The fourth case is symmetric to the second, and a good circuit is given by $ucy_2y_1z_{-1}ux_1x_2bcZ'z_{-1}Z''bu$ where we have split Z at z_{-1} to separate the segments contained in different faces.
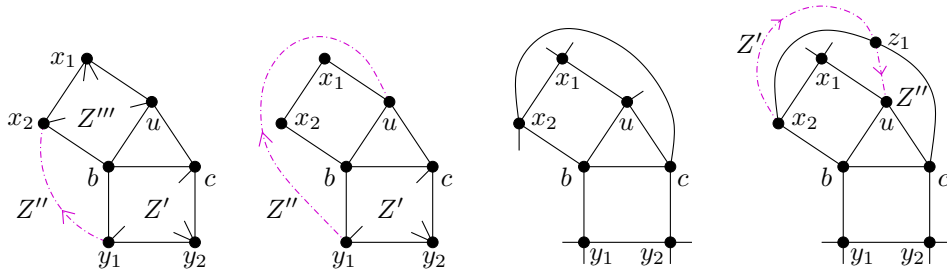

Figure 4.17: Configurations in case 1c with $X = x_1x_2$ and $Y = y_1y_2$.

Case 2: Resolving a 4-cycle of the form ubvcu

Since we are using restricted peggings, we may assume that bc is an edge in

G. This means that if we ignore the edges in the cycle *ubvcu*, the Eulerian circuit in the rest of the graph can be of the form $uXbcYvZu$, $uXvYbcZu$ or $uXvYcbZu$, assuming that we visit either *b* or *c* before *d* since orientation is unimportant. Furthermore, we note that $uXvYcbZu$ is just the first form in reverse while $uXvYbcZu$ can be obtained from the first by reversing the circuit and reflecting the whole picture in the vertical line through *u* after having drawn it as in Figure 4.18. That is, it is enough to assume the form $uXbcYvZu$. Again using the definition of pegging together with the argument from case 1a, we observe that *X* has at least 2 vertices. It turns out that there is an easy solution if *X* has at least 3 vertices, so we deal with this case first.
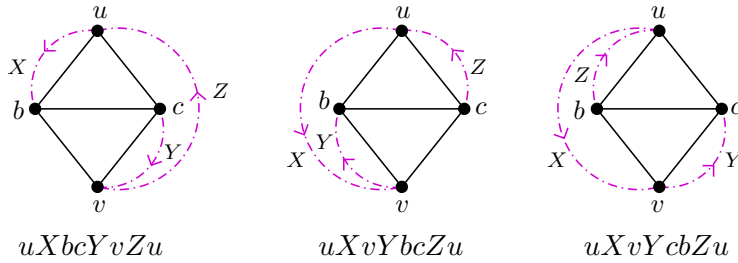

$$uXbcYvZu \qquad uXvYbcZu \qquad uXvYcbZu$$

Figure 4.18: Circuit configurations with a short subcycle of the form *ubvcu*.

Case 2a. Suppose that *X* has at least 3 vertices. Then $ubcYvZucvbX^{-1}u$ is 4-locally self-avoiding unless *Z* is a single edge and *Y* has only one vertex, say *y*. Under those assumptions though, *y* would be trapped (see Figure 4.19).
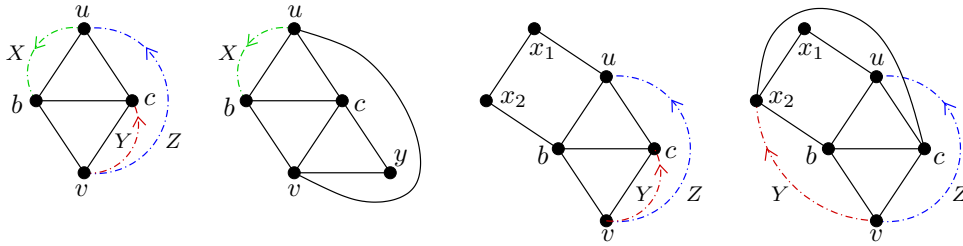

Figure 4.19: Cases 2a and 2bi. Green represents *X*, red *Y* and blue *Z*.

Case 2b. We now assume that $X = x_1 x_2$, and break into three subcases depending on the length of *Y*.

2bi. If *Y* has at least 3 vertices as in the third picture of Figure 4.19, consider the circuit $ubcux_1x_2bvcYvZu$. Since having $y_1 = x_2$ would obstruct *Z* (rightmost picture in Figure 4.19), we can be sure that x_2bvcy_1 is not a subcycle. Apart from *ubcu*, which we have created by design, there are no other potential local conflicts. Thus, we are done by case 1.

2bii. If $Y = y_1y_2$, then we first verify that these are distinct from x_1 and x_2.
Indeed, both $x_1 = y_1$ and $x_2 = y_2$ obstruct Z, whilst $x_1 = y_2$ and $x_1 = y_1$
both create 2-cuts. Hence, the configuration is that shown in Figure 4.20.
Let's try reroute to the circuit $ux_1x_2bcvZubvy_2y_1cu$. If $z_{-1} = y_2$ then we
would have a subcycle $z_{-1}ubvy_2$, but this traps either y_1 or x_1 and x_2.
The only other potential conflict occurs is x_2bcvz_1. If $z_1 = x_2$, then a 3-
cycle unslide centred at b can be applied; this requires firstly that $ux_2 \notin$
$E(G)$, which holds since the existence of that edge would either trap x_1
or imply that ab is in G, and secondly that c and v have only b as a
common neighbour, which also holds since $N(c) = \{u, b, v, y_1\}$ and $N(v) = \{b, c, y_2, z\}$ where $z \in G$ is distinct from all of the vertices labelled so far.
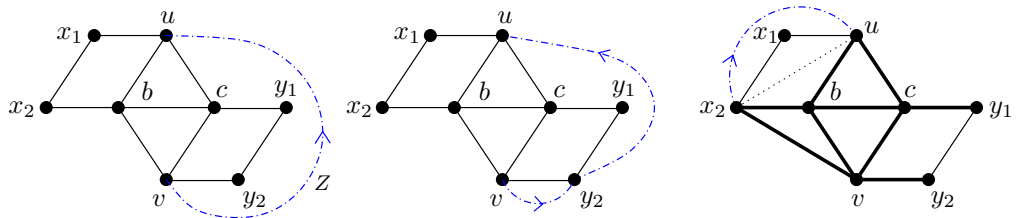

Figure 4.20: Configurations in case 2bii.

2biii. If Y has a single vertex y, then we have three adjacent triangles as shown in
bold in Figure 4.21 and we would like to apply a 3-cycle unslide centred at
c. There are two reasons why this may not be possible. The first is if there
is an edge uy, or equivalently $z_{-1} = y$. In this case, $ux_1x_2bcvZ'ycubvyu$
is 4-locally self-avoiding where $Z = Z'y$. Note that $cvZ'yc$ is not a short
subcycle since Z' contains at least x_1 and x_2. The second possibility is that
there is no edge uy, but $z_1 = x_2$. Let's split Z into $x_2Z'yZ''$, and consider
$ux_1x_2vcyZ''ubvyZ'^{-1}x_2bcu$. Since uy is not in G, it follows that Z'' has at
least one vertex so $yZ''ubvy$ is not a problem. However, Z' may be an edge.
If it is, we can avoid the subcycle $bvyx_2b$ by taking $ux_1x_2ycvbuZ''^{-1}yvx_2bcu$
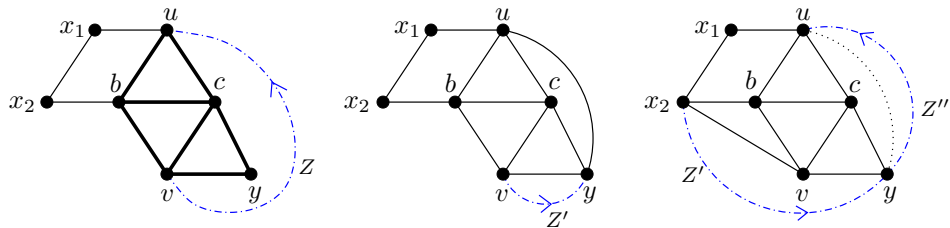which now has no possible conflicts.

Figure 4.21: Configurations in case 2biii.

Case 3: Resolving a 4-cycle of the form ubvdu

As in the previous cases, we split the induced circuit at each of the vertices on the problematic 4-cycle, name u, b, d and v, and consider all the possible orderings in which these are visited. In addition, we know that bc and uc are both in G, so we also keep track of where in the circuit these edges may be present. Ignoring the short subcycle, there are eight possible forms that the Eulerian circuit of $G - ubdvu$ may take, shown in Figure 4.22.
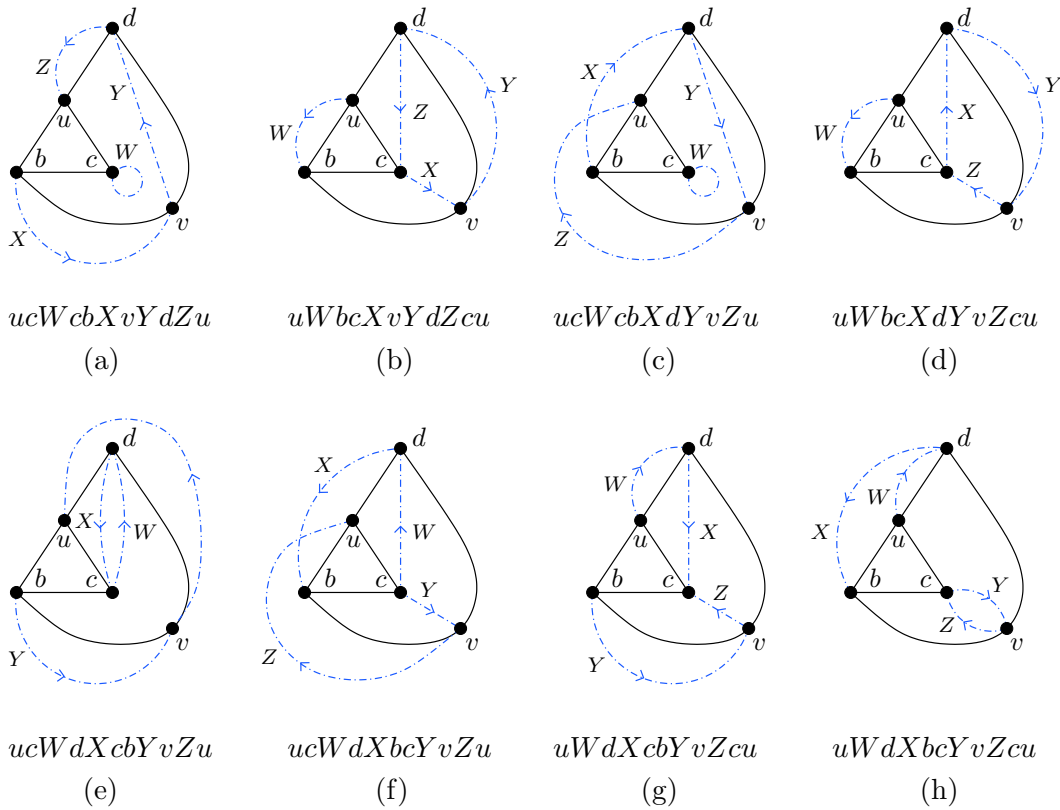

$$ucWcbXvYdZu$$

(a)

$$uWbcXvYdZcu$$

(b)

$$ucWcbXdYvZu$$

(c)

$$uWbcXdYvZcu$$

(d)

$$ucWdXcbYvZu$$

(e)

$$ucWdXbcYvZu$$

(f)

$$uWdXcbYvZcu$$

(g)

$$uWdXbcYvZcu$$

(h)

Figure 4.22: Induced Eulerian circuits of $G - E(ubvdu)$.

Although the embeddings shown in the figure are not necessarily the only possible ones, they do reflect two useful properties. First, the known cyclic ordering at u implies that the unique segment leading to u that does not contain ud, ub or uc must be in the region bounded by $ubvdu$ that does not contain c. Furthermore, a property of our chosen splitting is that each path segment is either contained inside $ucbvdu$ or outside $ubvdu$, noting that $ubcu$ is necessarily facial. When rerouting, it is important to keep in mind that this Eulerian circuit of $G - ubdvu$ is 4-locally self-avoiding as a walk but not as a circuit (i.e. not across u).

3a. Suppose $ucWcbXvYdZu$ is the induced 4-locally self-avoiding circuit. Then $ucWcvdubXvYdZu$ is a good Eulerian circuit since Z is disjoint from W, where the subpath $bXvYdZu$ is unchanged from the original circuit.

3b. An induced circuit of the form $uWbcXvYdZcu$ is the trickiest to resolve among these subcases. If Y has at least three vertices, then the circuit $ubcXvYdvbW^{-1}udZc$ has only one short subcycle (it is not possible that $y_{-1} = w_{-1}$, as this produces a 2-cut), namely $cubc$. We have $N(c) = \{u, b, x_1, z_{-1}\}$, and none of the neighbours of u and b can be inside the region bounded by $ducbv$. In particular, none of uz_{-1}, ux_1, bz_{-1}, bx_1 are in G. This shows that we may unpeg at c instead of at u, and the subcycle $cubc$ can now be resolved by case 1. Moreover, this is a restricted unpegging since $ub \in G$. The required configured is shown in bold in Figure 4.23.

Otherwise, we try the circuit $ucXvduWbvYdZcbu$. As long as this has no short subcycles except for $cbuc$, we can again unpeg at c and apply case 1. The possible problem is if $y_1 = w_{-1}$, creating a 4-cycle y_1dvby_1. This is cannot occur if $Y = y_1$, since $\{u, y_1\}$ would then be a 2-cut, so we may assume that $Y = y_1y_2$. In that case, let $W = W'y_1$ and reroute to $ubcZ^{-1}dy_2y_1bvduW'y_1vX^{-1}cu$. Now $cubc$ is the only short cycle.
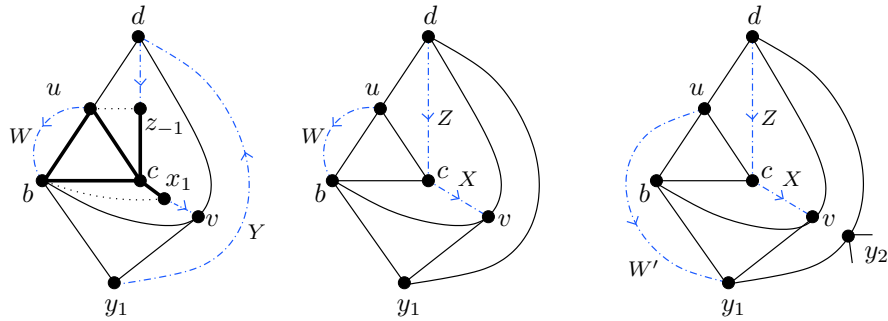

Figure 4.23: Configurations in case 3b.

3c. For $ucWcbXdYvZu$, a possible extension is given by $ubXdYvZudvbcWcu$. This is free of short subcycles as $bXdYvZu$ is copied directly from the original circuit, and W and X are disjoint, being contained in opposite regions bounded by the cycle $ucbvdu$.

3d. The form $uWbcXdYvZcu$ differs from that in part 3b only by the orientation of the segments. The same solution, after appropriate relabelling of the rerouted circuit, carries over.

3e. Suppose we have the form $ucWdXcbYvZu$. We know that Z is in the outside region bounded by $ubvdu$, whilst W and X are inside. If Y is also inside then $\{u, d\}$ would be a 2-cut, and if Y is outside then $\{c, d\}$ would be a 2-cut.

3f. A circuit of the form $ucWdXbcYvZu$ can be rerouted to $ucYvZudvbcWdXbu$. There is only place a short subcycle may occur, which is at $bcWdXb$ in the situation that W and X have only 3 edges in total. However, we know that a lies on X, which therefore has at least two edges, and cd is not in G so W also has at least two edges.

3g. For $uWdXcbYvZcu$, we reroute to $uWdvbcX^{-1}dubYvZcu$. The subpath $y_{-1}vZcuw_1$ cannot be a short subcycle since having Z be an edge and $y_{-1} = w_1$ produces a 2-cut $\{b, w_1\}$. There are no other potential conflicts.

3h. Finally, the induced circuit $uWdXbcYvZcu$ leads to the same contradiction as we saw in 3e.

$\square$

4.3.3   Proof of the main theorem

All of the hard work we have done reduces the proof of the main theorem to a straightforward induction argument.

Proof of Theorem 4.17. To establish the base case, we exhibit Eulerian circuits that are 4-locally self-avoiding for each of the antiprisms except the octahedron. We must also do this for each graph obtained from the octahedron by one expansion operation, but there is only one such graph. To see this, note that the octahedron does not have the local configuration required for pegging. Any 3-cycle slide results in a tfp-graph on 8 vertices, but the 4-antiprism is the unique such graph. Finally, 4-cycle addition results in a tfp-graph on 10 vertices. There are three isomorphism classes of such graphs shown in Appendix A – this, we know from the generation results in Figure 3.16 – of which one is the 5-antiprism, and one can be obtained from the 4-antiprism by pegging twice. Consequently, it turns out that there is only one additional graph to check here. This and the antiprisms are shown in Figure 4.24, and have the following 4-locally self-avoiding Eulerian circuits with vertex labelling given in the figure;

- the 4-antiprism admits the circuit $abcy_2y_1x_1x_2x_3cax_1by_1x_2y_2x_3a$,

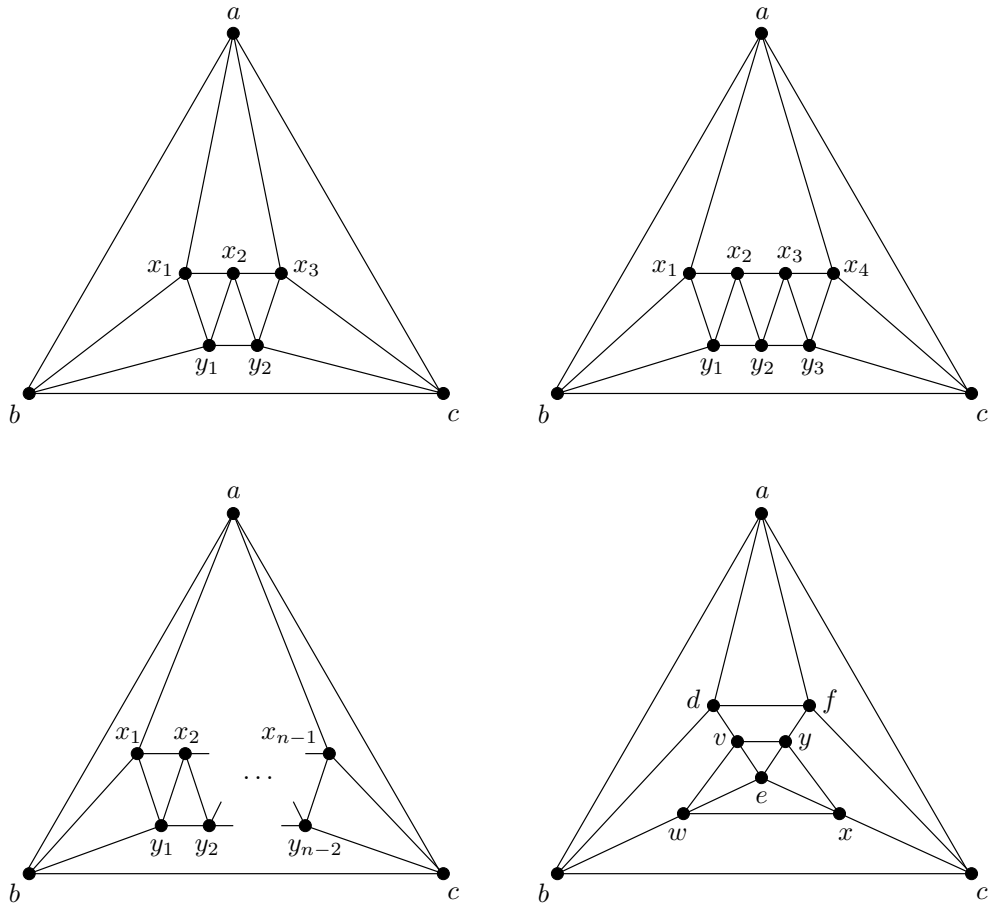- the 5-antiprism admits $abcy_3y_2y_1x_1x_2x_3x_4cax_1by_1x_2y_2x_3y_3x_4a$,

Figure 4.24: Graphs isomorphic to the 4-antiprism, the 5-antiprism, the n-antiprism, and the octahedron with a single 4-cycle addition.

- following the previous two particular cases, for a general n-antiprism we may take $abcy_{n-2}\ldots y_2y_1x_1x_2\ldots x_{n-1}cax_1by_1x_2y_2x_3\ldots x_{n-2}y_{n-2}x_{n-1}a$,

- the graph obtained from the octahedron by applying a single 4-cycle addition admits $abwevyxcbdafyexwvdfca$.

Given an arbitrary tfp-graph G, Corollary 4.18 implies the existence of a sequence of tfp-graphs $G_0, G_1 \ldots G_n$ such that $G_n = G$, G_0 is one of the graphs mentioned in the base case, and G_i can be constructed from G_{i-1} by applying either a pegging operation, a 4-cycle addition or a 3-cycle slide for all $1 \leq i \leq n$. We have verified above that G_0 has a 4-locally self-avoiding Eulerian circuit. Lemmas 4.20, 4.21 and 4.26 imply that each G_i in the sequence also admits such an circuit. In particular, this includes G. □

Our result can also be weakened to one concerning path decompositions.

Corollary 4.27. *A 3-connected quartic planar graph G has a P_5-decomposition if and only if it has even order. In addition, if G has odd order, then it has a near-P_5-decomposition with two edges left over which can be chosen to be adjacent or not.*

Proof. If G is isomorphic to the octahedron, we may take the decomposition shown in Figure 4.25. For all other 3-connected quartic planar graphs, a suitable P_5-decomposition can be obtained by taking a 4-locally self-avoiding Eulerian circuit from Theorem 4.17 and, starting from an arbitrary vertex, cutting the circuit into paths consisting of the next four consecutive edges. This requires the number of edges to be divisible by four, which holds by the assumption that the graph has even order. Conversely, if G is a tfp-graph with a P_5-decomposition, then $|E(G)| = 2|V(G)|$ must be divisible by four, and hence G has even order.

If G has odd order, then $|E(G)| \bmod 4 = 2$ so the same cutting process will yield paths of length 4 that leave two edges uncovered. To make these adjacent, we can simply take the two edges left over at the end. To make them non-adjacent, at two points in the cutting procedure we can skip a vertex and then take the next four consecutive edges. $\qquad\square$

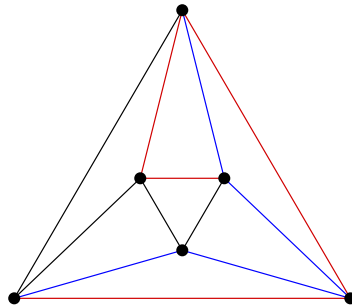

Figure 4.25: A P_5-decomposition of the octahedron.

4.4   Open problems

When it comes to determining the actual value of k, we have just scratched the surface and open problems in this area abound.

For the particular case we have studied, the obvious strategy moving forward is to see whether the rerouting method can be pushed further. In attempting this, we see two main complications. For one, the amount of casework will increase because there are more local conflicts that must be checked and resolved. The

more subtle issue is that the local conflicts were previously 3-cycles or 4-cycles, and we knew enough about the cyclic orderings at two or three of those vertices to be able to conclude in many instances that some vertex would be trapped. The most difficult cases involved 4-cycles for which two of the vertices were neither part of the initial problem cycle nor the endpoints of path segments, since in this case it is possible for a segment to enter at one vertex and leave at the other, leaving nothing trapped. If we move up even just to to 5-cycles, this situation arises much more often. It may be possible to overcome this issue by splitting segments at more vertices and finding alternative routes instead of deriving contradictions, but again this increases the casework significantly.

Although our question included a clause to allow finitely many counterexamples, we only had one exception, and it would be informative to know how quickly this increases as we try to avoid longer cycles. The computational results summarised in the following table show the number of quartic planar graphs of each realisable order up to 20 that do not have ℓ-locally self-avoiding Eulerian circuits for $4 \leq \ell \leq 8$. Blank entries represent zeros.

ℓ	6	8	9	10	11	12	13	14	15	16	17	18	19	20
4	1													
5	1	1												
6	1	1	1											
7	1	1	1	3		1								
8	1	1	1	3	2	7	3	5						

The first row supports our result, and on the whole, there are not very many small exceptions. It is tempting to guess that that our upper bound is sharp and that the exceptions counted above are the only ones, since size was the only obstruction in the 4-locally self-avoiding case. No gadgets forcing smaller subcycles have been identified in computer searches undertaken so far, although we cannot exclude this possibility.

A related question that is far more attainable is whether we can characterise the obstructions to having a 4-locally self-avoiding Eulerian circuit in the toy version we considered. There is a lemma in [1] which, loosely speaking, says that if an Eulerian graph with maximum degree 4 has a cutset S satisfying some mild additional conditions such that the components of $G - S$ has triangle-free Eulerian circuits, then these circuits can be stitched together to obtain a triangle-free Eulerian circuit of G.

A similar idea may be workable for 4-locally self avoiding circuits in quartic planar graphs. Suppose we could decompose any quartic planar graph into 3-connected blocks that satisfy the conditions of Theorem 4.17. Moreover, suppose that given good Eulerian circuits in each blocks, they can be stitched together into a good Eulerian circuit of G. Then if a graph does not admit a good Eulerian circuit, at least one of its blocks must be isomorphic to the octahedron by Theorem 4.17. Observe that the diamond attachment is precisely obtained by cutting one edge of the octahedron. This means that an affirmative answer to the following is plausible.

Question 4.28. Is the diamond attachment the only obstruction for a quartic planar graph to have a 4-locally self-avoiding Eulerian circuit?

We conjecture that it is, supported by a computer check up to 20 vertices.

Chapter 5

Circle representations

Given a collection of circles drawn on a plane, we can obtain a planar quartic graph by taking as vertices the kissing and crossing points between the circles, and letting the edge set be given by the circular arcs between those points. This system of circles is then said to form a *circle representation* of the graph. Lovász [56] conjectured that every simple connected quartic planar graph has such a representation. However, an infinite family of counterexamples were given by Bekos and Raftopoulou [9], who also proved that 3-connectedness is sufficient to guarantee the existence of a circle representation. In this section, we investigate smaller counterexamples to the conjecture, and show that the 3-connected condition can be weakened a little.

5.1   Geometric graphs and graph drawing

We have so far considered abstract graphs as well as embedded graphs. The next step in this progression of increasing structural rigidity is to view graphs as geometric objects drawn in the plane (or possibly on other surfaces). This is the view taken in the area of geometric graph theory, under which our third problem falls. A basic question here is to ask what kind of graphs can be drawn in some special, predefined way involving geometric objects. A standard example to illustrate this is Fáry's theorem:

Theorem 5.1 (Fáry [33])**.** *Every planar graph can be drawn with vertices represented by points on the plane and edges as straight line segments.*

Results of this sort in geometric graph theory are closely related to graph drawing, which is concerned with constructing geometric representations of graphs

that are geared toward satisfying some specific aesthetic criteria. These considerations might include symmetry, minimal area, or number of bends. As much as we like pretty pictures, the motivation for this study comes primarily from applications. For example, those qualities we listed are important in graph visualisation, VLSI circuit design, and train networks respectively. We will mention a few more applications of specific theorems in the next section.

Our focus is on circle representations of graphs, which are best introduced in the context of some more classical types of graph representations that we now briefly survey.

5.1.1   Contacts, coins and circles

Broadly speaking, a *contact representation* of a planar graph is a drawing of a graph on the plane such that vertices correspond to geometric objects, and edges are defined by contacts between objects. Perhaps the most classical case is that taking the objects to be disks, with edges corresponding to two disks that are tangent. Specifically:

Definition 5.2. A *coin representation* of a graph G is a collection of interior-disjoint circles in $\mathbb{R}^2$ such that the circles are in bijective correspondence with $V(G)$, and two vertices are adjacent in G if and only if their corresponding circles are tangent to each other. Any graph that can be represented in this way is called a *coin graph*.

Given some such collection of circles, it is clear that the graph for which it is a coin representation must be planar. The remarkable fact that all simple planar graphs admit such a representation is known as the circle packing theorem, and was originally proved by Koebe in 1936.

Theorem 5.3 (Koebe [53])**.** *Every simple planar graph is a coin graph.*

Several independent proofs of this result have been published. For a survey of these, we refer to [74] which also discusses some fascinating interconnections between circle packing and other areas of mathematics. For example, applications of circle packing include proofs of the planar separator theorem in graph theory (see [60]) and the Koebe Uniformisation Theorem in complex analysis (see [44]).

The proof of the circle packing theorem given by Brightwell and Scheinerman [14] actually gives an even stronger type of representation called a *double circle representation*. This consists of vertex-circles $\{C_i\}$ and face-circles $\{D_j\}$

such that the vertex-circles form a circle representation of G, the face-circles form a circle representation of G^* except that the circle corresponding to the outer face contains all of the others, for each $xy \in E(G)$ there is a point (called an edge-point) in the plane where the four circles corresponding to x, y, and the two faces incident to xy all meet, and a face-circle intersects a vertex-circle if and only if the vertex is incident to the face in G. For good measure, one can also throw in the property that at every edge-point the two vertex-circles cross the two face-circles at right angles. The existence of such a representation seems rather miraculous, and yet we have the following.

Theorem 5.4 (Brightwell and Scheinerman [14])**.** *Every 3-connected simple planar graph has a double circle representation.*
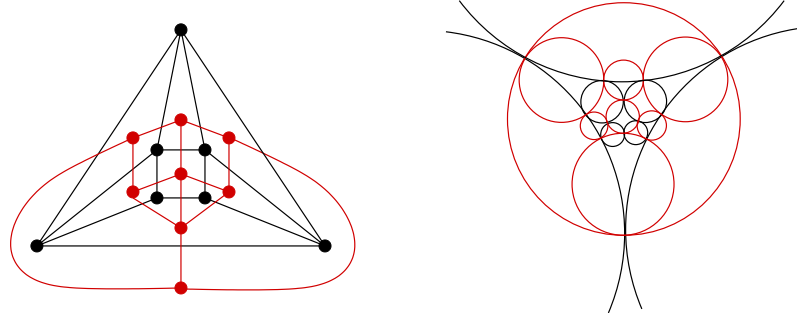

Figure 5.1: A graph and its double circle representation with the dual graph and face-circles shown in red.

Since the circle packing theorem, there has been great interest in characterising the planar graphs that can be represented as contact graphs by a range of geometric objects including arcs and line segments [2, 45], triangles [25, 38], squares and rectangles [34, 71, 76] and hexagons [29].

Another kind of graph representation is as an *intersection graph* of some geometric objects, which is more general than a contact representation where vertices correspond to geometric objects that are no longer required to have disjoint interiors, and edges correspond to the intersections of the objects. The circle representations that we will be studying are neither contact representations nor intersection representations in the classical sense, but they are closely related to the former and share some obvious properties with the latter. The main difference is that our vertices will be given by contacts while the edges are geometric objects, namely circular arcs that form full circles. It is clear that any graph obtained with these correspondences must be planar and quartic.

Definition 5.5. A *circle representation* of a quartic planar graph G is a collection of circles in $\mathbb{R}^2$ such that the kissing and crossing points of the circles are in bijective correspondence with $V(G)$, and the circular arcs between those points correspond to the edges of G.

Here, we borrow the term *kissing point* from geometry to refer to the point of intersection between two tangent circles, whilst the self-evident *crossing point* refers to a point of intersection between two circles whose interiors are not disjoint. An example of a quartic planar graph and its circle representation is shown in Figure 5.2.
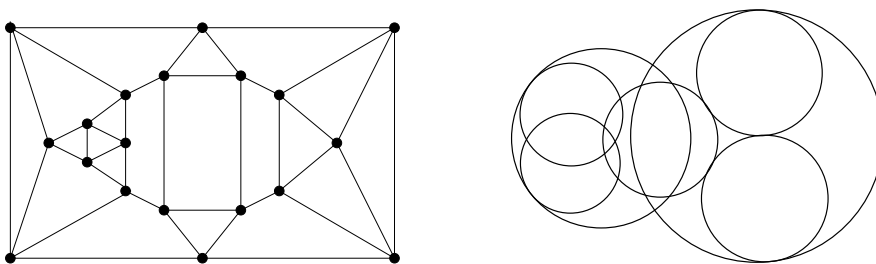

Figure 5.2: A circle representation of a simple quartic graph.

There is a close connection between coin representations and circle representations which is encoded in the following relationship.

Lemma 5.6. *For any quartic graph G, a coin representation of TG is also a circle representation of G.*

Proof. In a coin representation of TG, two circles are tangent if and only if the vertices they represent are adjacent in TG if and only if the corresponding two faces of G have a vertex in common. Since our graph is quartic, every vertex is shared by precisely two coloured faces in a chessboard colouring of G, and hence represented by a kissing point. In addition, each edge xy in G is incident to precisely one coloured face which corresponds to a vertex in TG and hence a circle in G. The kissing points along the circle correspond to vertices along the cycle bounding this coloured face in G. In particular, there is an arc between the points representing x and y that gives the required representation of xy. $\square$

The reader may find it illustrative to trace out the chain of correspondences in the above proof while looking at the example shown in Figure 5.3. Observe that due to the application of the circle packing theorem, the construction in the proof produces circle representations that only represent vertices as kissing point,

and also in which all circles have disjoint interiors. This is a point that we will return to later.
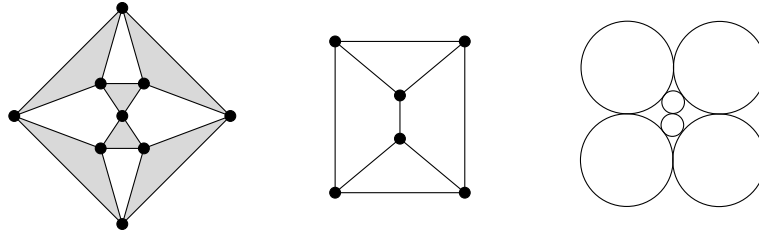

Figure 5.3: A coin representation of TG is a circle representation of G.

It is worth mentioning that circle representations are not the only type of representation specific to quartic planar graphs that has been considered. For example, orthogonal drawings of quartic planar graphs in which each vertex corresponds to a point on a grid and each edge is a sequence of gridline segments have been widely studied in relation to VLSI circuit design and floorplanning (see for instance [57]). Our particular interest in circle representations is inspired by the following conjecture of Lovász published in 1970, which is analogous to the circle packing theorem.

Conjecture 5.7 (Lovász [23, 31])**.** Every simple quartic planar graph admits a circle representation.

This is already known to be false, however the obstructions to circle representability are not well understood. Our goal is to work toward a characterisation of the circle representable graphs which we approach from two directions by, on one hand, finding new counterexamples, and on the other, finding more classes of graphs that do admit circle representations.

5.2   Counterexamples to Lovász' conjecture

In 2012, Bekos and Raftopoulou published the first known counterexamples of Conjecture 5.7 and in fact gave an infinite family of them, followed by a second family of 2-connected counterexamples in an update in 2015 (see [9]). Among those graphs, the smallest had 822 vertices. We will record those constructions, and then devote the remainder of this section to presenting two new infinite families of counterexamples, the smallest of which have order 68.

5.2.1   Two known infinite families

The counterexamples gives in [9] are defined by attaching certain gadgets to the graph of the octahedron. We start with the subdivision of the octahedron which consist of the vertices and solid edges in Figure 5.4(c). Each dotted edge in that diagram indicates a place where we will attach one copy of the gadget graph shown in Figure 5.4(a). To do this, we simply identify the vertices of degree 2 in the gadget with the vertices of degree 2 at either end of a dotted line in the main octahedron graph so that the two gadget edges are consecutive in the cyclic ordering at each identified vertex, as are the two octahedron edges. This rotation system is the obvious one indicated in the picture that we get by scaling and rotating the gadget as necessary so that the lower circular arc in the gadget matches up with a dotted circular arc. The resulting graph has 96 cutvertices, which come from the 48 copies of the gadget.

A 2-connected counterexample can be obtained by replacing gadget (a) by that shown in Figure 5.4(b) in the above description. Both of the counterexamples described so far are on 822 vertices.
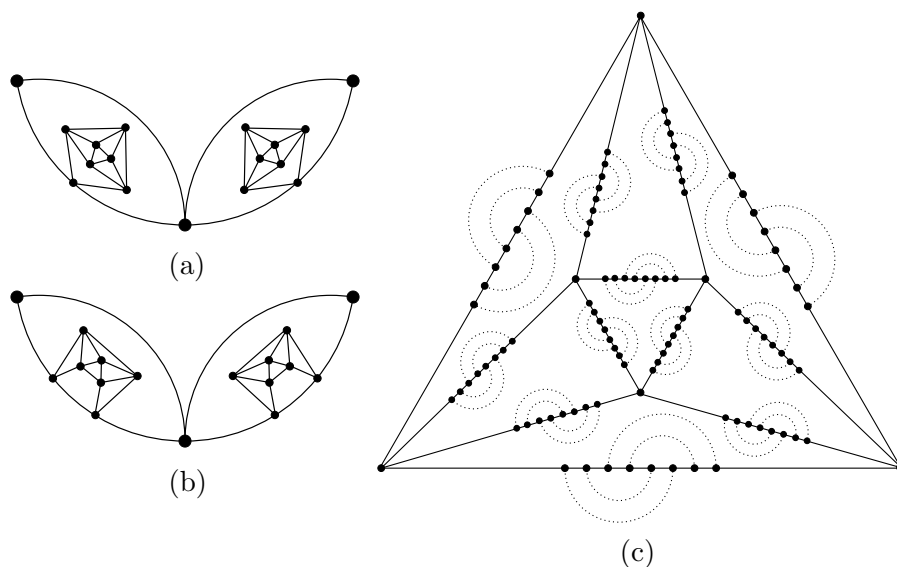

Figure 5.4: Pieces of the counterexample

The key step of the proof in [9] that these graphs are not circle representable is to show that if either one does have a circle representation, then the representation of each gadget is independent of the representation of the main octahedron. That is, each circle in the representation either represents edges that are all in a gadget or all in the octahedron. After some geometry, one finds that the con-
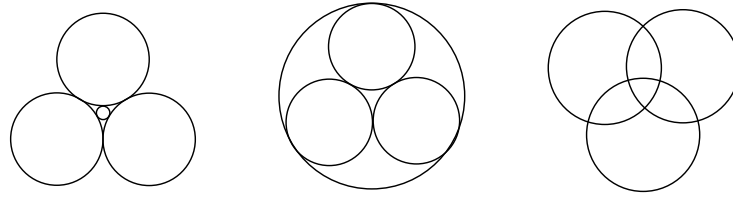
Figure 5.5: Three circle representations of the octahedron.

figuration forces at least one edge of the innermost face in the octahedron to be realised as a major arc of a circle. To reach a contradiction, it is proved that the octahedron has three different circle representations up to some technical notion of equivalence (these are shown in Figure 5.5), none of which have an arc satisfying the required property. We will not reproduce this argument in full since our own counterexamples are easier to prove, and the method we use is independent of their lemmas although we do use the idea of independent circles.

There are two ways to obtain infinite families of counterexamples from these counterexamples. The first is to modify the gadgets. Observe that gadget (a) consists of two copies of a 7-vertex subdivision of the octahedron, which is chosen purely because it is the smallest quartic planar graph. Bekos and Raftopoulou call each of these a *loop subgraph*. In general, any quartic planar graph with one edge on its outer face subdivided can be used as a loop subgraph to replace these. Similarly, gadget (b) has two *biloop subgraphs* which are again derived from the octahedron, but in general any 2-connected planar graph that has two vertices of degree two on its outer face and all other vertices are of degree 4 will do.

The other way to obtain an infinite family is to attach more copies of gadgets. The current construction uses four copies of the gadgets attached in pairs. The proof outlined above still goes through if one attaches more than two pairs – simply copy the existing configuration and repeat it as many times as we like along each edge.

5.2.2   A base multigraph

Although Lovász' conjecture was specific to simple graphs, the definition of a circle representation extends without change to multigraphs. Working with multigraphs is actually somewhat easier because there are a very limited number of ways in which a loop or digon can be represented. In this section, we give a 12-vertex multigraph that does not have a circle representation from which we construct simple quartic planar graphs with the same property.

Theorem 5.8. *The multigraph H shown in Figure 5.6(a) does not have a circle representation.*

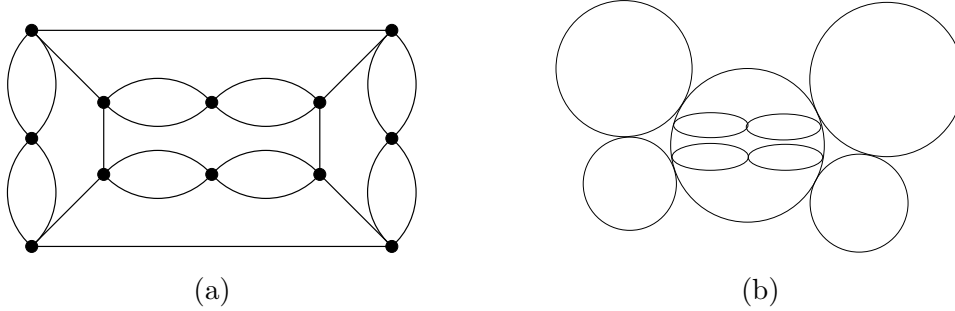<center>(a)   (b)</center>

<center>Figure 5.6: A non-simple counterexample</center>

The idea of the proof is to first show that any circle representation of this graph must have a particular structure forced by the neighbouring digons, indicated in Figure 5.6(b). This corresponds to a particular configuration of circles in the plane with restrictions imposed by certain pairs that must be kissing circles. The problem is then reduced to showing geometrically that this configuration is not realisable. We will handle this geometric aspect first.

Lemma 5.9. *Suppose we have four circles $\{C_i(r_i, t_i)\}_{i=1,2,3,4}$ in the plane with radii $r_i > 0$ and which touch the x-axis at points $(t_i, 0)$ respectively, and assume the circles are numbered so that $t_1 < t_2 < t_3 < t_4$. In addition, suppose that C_1 is tangent to C_2 and C_4, and C_3 is also tangent to C_2 and C_4. Let $N = t_4 - t_1$, $M = t_3 - t_2$, $L = t_2 - t_1$ and $R = t_4 - t_3$. Then*

(i) $MN = LR$, and

(ii) the lengths L and R are in $\{f_N^+(M), f_N^-(M)\}$ (not necessarily in that order)

where $f_N^+(M) = \frac{N-M+\sqrt{(M-N)^2-4MN}}{2}$ and $f_N^-(M) = \frac{N-M-\sqrt{(M-N)^2-4MN}}{2}$ for some fixed N and $0 < M \leq (3 - 2\sqrt{2})N$. In addition,

(iii) $f_N^-(M)$ is increasing, and

(iv) $f_{N_1}^-(M) > f_{N_2}^-(M)$ whenever $0 < N_1 < N_2$ and $0 < M \leq (3 - 2\sqrt{2})N_1$.

Proof. From Figure 5.7, we observe that

$$(t_4 - t_1)^2(t_3 - t_2)^2 = (4r_1r_4)(4r_2r_3) = (4r_1r_2)(4r_3r_4) = (t_2 - t_1)^2(t_4 - t_3)^2$$
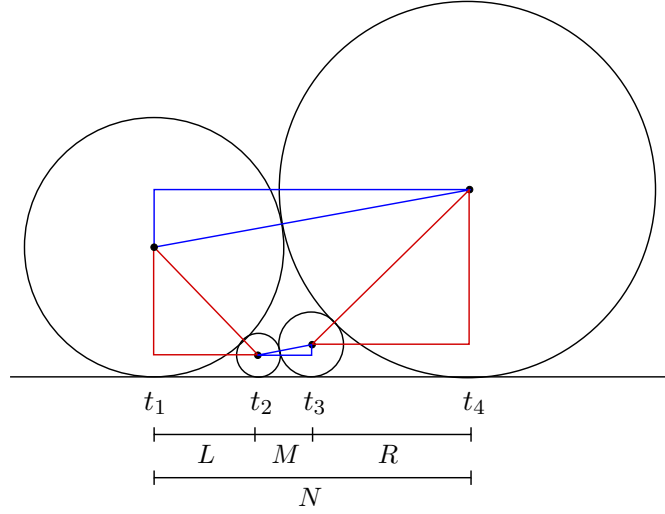
which implies (i).

Figure 5.7: Configuration of circles in Lemma 5.9.

For (ii) we can write $L = 1 - M - R$, and so from the first part we have

$$(N - M - R)R = MN$$
$$R^2 + (M - N)R + MN = 0$$
$$R = \frac{N - M \pm \sqrt{(M - N)^2 - 4MN}}{2}.$$

One of the solutions is R whilst the other is L, noting that we could equivalently have used the substitution $R = 1 - L - M$.

If we let $f(N, M) := f_N^-(M)$, then

$$\frac{\partial}{\partial M} f = -\frac{1}{2} \left(\frac{2(M - N) - 4N}{2\sqrt{(M - N)^2 - 4MN}} + 1 \right)$$

and the inequality $\frac{\partial}{\partial M} f > 0$ has as a solution $M < (3 - 2\sqrt{2})N$ and $N > 0$. This implies (iii). Similarly, one can verify that $\frac{\partial}{\partial N} f < 0$ on the specified interval to establish (iv). $\qquad \square$

Note that the circles described in the previous lemma are not necessarily disjoint. In particular, we do not exclude the possibility that C_1 and C_3 or C_2 and C_4 cross. These basic results allow us to show the impossibility of the following circle configuration. Take two sets of circles, $\{C_1, C_4, C_5, C_8\}$ and $\{C_2, C_3, C_6, C_7\}$, satisfying the conditions of the previous lemma so that we have eight circles in total, all tangent to the x-axis and within the sets some circles are tangent to each other as described above. Assume additionally that there are no tangencies
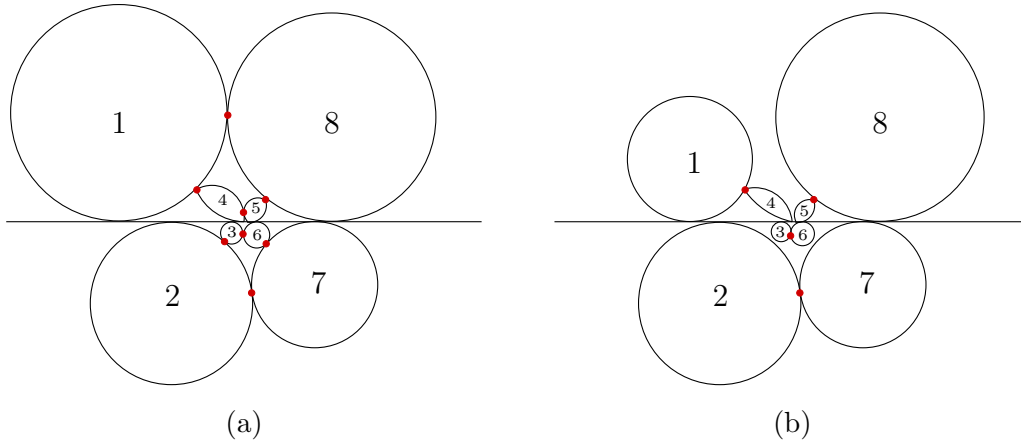
Figure 5.8: Impossible circle configurations.

between circles of different sets. It is convenient to view one set of four circles, say $\{C_1, C_4, C_5, C_8\}$ as being above the x-axis, and the other below. Moreover, adopting the notation for radii and tangent points with the axis defined before, we stipulate that $t_1 < t_2 < \ldots < t_8$. The tangencies are illustrated in Figure 5.8a in which all numbered simple closed curves should be interpreted as circles, tangent points between circles are marked in red, and the order of the tangent points with the axis may be inferred from the numbering.

Lemma 5.10. *The configuration described is impossible.*

Proof. By scaling and translating horizontally, we may assume that $t_2 = 0$ and $t_7 = 1$. Let $M = t_6 - t_3$, $M' = t_5 - t_4$ and $N' = t_8 - t_1$. Since $t_3 < t_4 < t_5 < t_6$ we know that $M' < M$. Also, as $t_1 < 0$ and $t_8 = t_1 + N' > 1$ we must have $N' > 1$. Lemma 5.9(ii) tells us that the lengths $t_4 - t_1$ and $t_8 - t_5$ are given by $\frac{1}{2}(N' - M' \pm \sqrt{(M' - N')^2 - 4M'N'})$, so using our earlier notation the length of the shorter interval is $f_{N'}^-(M')$. Applying Lemma 5.9(iv) then (iii), we find that

$$f_{N'}^-(M') < f_1^-(M') < f_1^-(M).$$

On the other hand, we know that $f_1^-(M) = \min(t_3 - t_2, t_7 - t_6)$, so $f_{N'}^-(M') < t_3 - t_2$ and $f_{N'}^-(M') < t_7 - t_6$. Now if $f_{N'}^-(M') = t_4 - t_1$, then we would have

$$t_4 = t_1 + (t_4 - t_1) < 0 + f_{N'}^-(M') < t_3 - t_2 = t_3$$

which is a contradiction. Similarly, if $f_{N'}^-(M') = t_8 - t_5$ then

$$t_5 = t_8 - (t_8 - t_5) > 1 - f_{N'}^-(M') > 1 - (t_7 - t_6) = t_6$$

so we reach a contradiction in both cases.

□

Corollary 5.11. *The configuration shown in Figure 5.8b cannot be realised by circles.*

Proof. It is enough to show that it is possible to adjust the circles to create new touching points while preserving their order along the axis, so that we obtain the configuration in the preceding lemma. Below the axis, fix C_2 and C_6 and replace C_3 and C_7 with two new circles C_3' and C_7' that are both tangent to C_2, C_6 and the line, as shown in Figure 5.9. Then $r_3' > r_3$, meaning $t_6 - t_3 = 2\sqrt{r_3 r_6} < 2\sqrt{r_3' r_6} < t_6 - t_3'$, so $t_3' < t_3 < t_4$. Also, we know that $t_2 < t_3'$ by choice of C_3' being tangent to C_2 and C_6, so this replacement preserves the order of the circles. Similarly, we have $r_7' < r_7$ from which we deduce that $t_6 < t_7' < t_7 < t_8$. Since the inequalities are strict at each step, no circle above the axis is tangent to any circle below the axis. Doing the same thing above the axis by fixing C_4 and C_8 and replacing C_1 and C_5 produces the desired configuration, and then we are done by Lemma 5.10. □
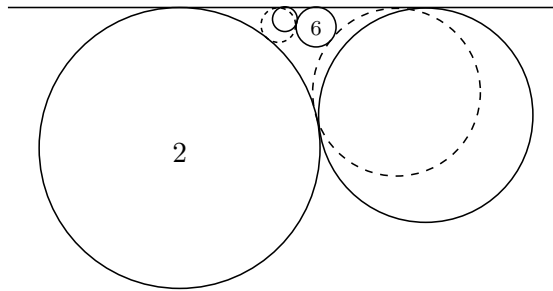

Figure 5.9: Replacing C_3 and C_7 by the circles C_3' and C_7' with dashed edges.

Proof of Theorem 5.8. We begin by making two observations. Firstly, any pair of digons sharing exactly one vertex must be realised by two circles that are tangent at that shared vertex. This is because if one of the digons is produced by crossing circles, then the edges of the neighbouring digon are realised by arcs of the same two circles which is only possible if the two digons share both of their vertices. We also observe that for a quartic planar graph, if a vertex v of a digon is not a cut-vertex, then in any embedding of the graph the parallel edges must be consecutive in the cyclic ordering at v.

The first observation implies that any circle representation of H must have one circle representing each digon. One of the circles on which v_1 lies therefore corresponds to a digon, so the remaining edges e_1 and e_8 must lie on the same circle. Furthermore, by the second observation it follows that v_1 is a tangent

point. The same argument applies at each vertex v_i for $i = 1, 2, \ldots, 8$ so we find that $e_1, e_2, \ldots, e_8$ all lie on the same circle and hence the cycle $e_1, e_2 \ldots, e_8$ must be represented by a single circle. In addition, all vertices have now been shown to be tangent points, so no circles cross. Now if a circle represetnation of H exists, it must have one circle C_0 corresponding to the 8-cycle, and then 8 more circles $C_1, \ldots, C_8$ labelled so that v_i is the tangent point of C_i with C_0, and these points occur in the cyclic order around C_0. In addition, (C_1, C_4), (C_2, C_3), (C_5, C_8) and (C_6, C_7) are pairs of kissing circles, and no other circles kiss.

These tangency relationships are the same as those in the configuration in Figure 5.8b if we take C_0 to be the line. Indeed, if we view the circle representation on the sphere, rotate so that a point on C_0 strictly between v_1 and v_8 is the North Pole and then apply a stereographic projection, the resulting planar drawing satisfies the conditions of Corollary 5.11. As that configuration was not realisable by circles, no circle representation of H can exist. $\square$

5.2.3   Small simple counterexamples

With our multigraph counterexample in hand, we can construct simple counterexamples by taking a subdivision to obtain a simple graph, and then planting gadget attachments at the degree 2 vertices to fix it up into a quartic graph. Our smallest counterexamples are obtained by taking the smallest possible attachments, shown in Figure 5.10, which were also used in the previous counterexamples. Here, the solid lines and vertices are *gadget-vertices* and *gadget-edges*, whilst the broken lines and uncoloured vertices will be referred to as *non-gadget*. Let the *vertices of attachment* of a gadget be those gadget vertices that are adjacent to non-gadget vertices. That these are the smallest possible follows from the fact that they are based on the octahedron graph; gadget (a) is a subdivision of the octahedron, while the octahedron can be recovered from gadget (b) by identifying the two attachment vertices.

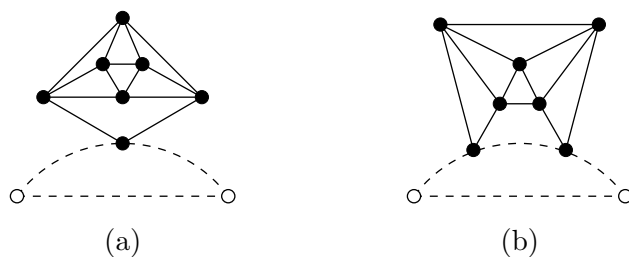(a)   (b)

Figure 5.10: Gadget subgraphs

In a circle representation, every edge is an arc of exactly one circle, so the edges along a circle represents a simple cycle in the graph. Let's call such a cycle *circular*, and let C_{xy} denote the unique circle that contains the edge xy. Geometrically, one should think of C_{xy} as an honest circle, but notationally we will use it to denote the edge set of the circular cycle. If a gadget has exactly one vertex of attachment a, then a circular cycle that contains both gadget and non-gadget edges would have to contain a twice (in order to 'enter' and 'exit' the gadget, if we orient the cycle) which is impossible. By similar reasoning, if a gadget has exactly two vertices of attachment then any circular cycle containing both gadget and non-gadget edges must enter the gadget at one vertex a and leave at the other, say b, in order to avoid subcycles. That is, both vertices of attachment lie on the circle. We have thus established the following lemma.

Lemma 5.12. *Suppose we have a circle representation of a graph containing a gadget subgraph H.*

1. *If H has exactly one vertex of attachment, then every circle in the representation contains either gadget-edges only, or non-gadget edges only.*

2. *If H has exactly two vertices of attachment, then every circle containing both gadget and non-gadget edges must contain both vertices of attachment.*

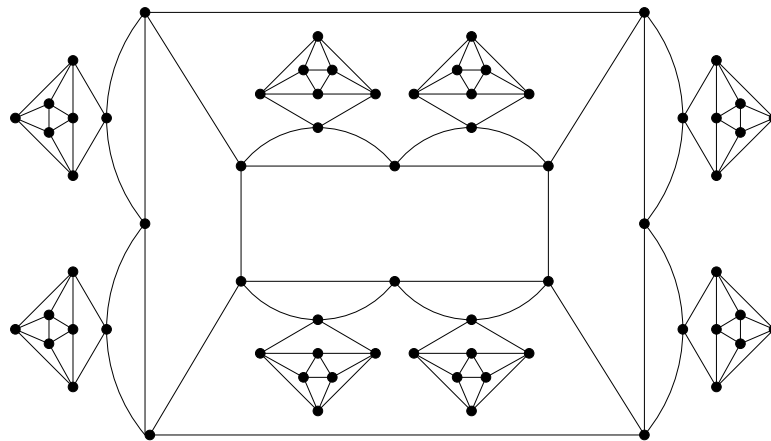**Theorem 5.13.** *The graph shown in Figure 5.11 does not have a circle representation.*

Figure 5.11: A counterexample on 68 vertices.

Proof. Suppose we have a circle representation of G, and let $H \subset G$ be one of the gadgets. Since each gadget has a single vertex of attachment, it follows from

Lemma 5.12 that every circle contains either gadget edges only, or non-gadget edges only. That is, we have a circle representation of H that is independent from the rest of the graph. In particular, if we delete the circles containing only gadget edges, we are left with a circle representation of $G - H$ but with an additional parallel edge between the neighbours of the vertex of attachment.

An identical argument applied to the resulting multigraph shows that each of the remaining gadgets in the graph are realised independently by sets of circles that only contain edges and vertices of that particular attachment. Deleting each of those circles from the original representation of G gives a circle representation of H, but this contradicts Theorem 5.8.   □

We can obtain a 2-connected counterexample by replacing each gadget above with one that has 2 vertices of attachment. The argument is again by contradiction and uses the same idea as before, which relies on showing that the gadgets can be detached to give a circle representation of the same base multigraph.

Theorem 5.14. *The 2-connected graph shown in Figure 5.21 does not have a circle representation.*
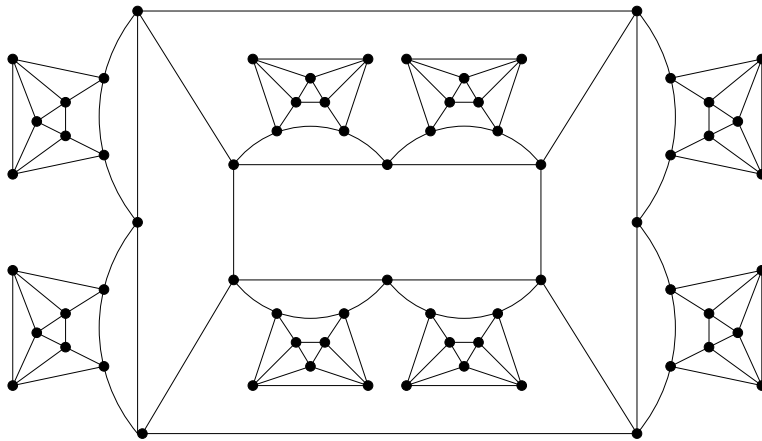

Figure 5.12: A 2-connected counterexample on 68 vertices.

Proof. Suppose we have a circle representation of G and consider the subgraph induced by one of the gadgets, say H_{ab}, together with the neighbours of its attaching vertices, labelled as shown in Figure 5.13. First suppose that $C_{va} \neq C_{bw}$. If ab lies on C_{va}, then C_{va} must contain both gadget and non-gadget edges. If we traverse the cycle of edges of C_{va} so that we enter the gadget at b, then in order to return to vertex v we are forced to leave the gadget at either a or b,
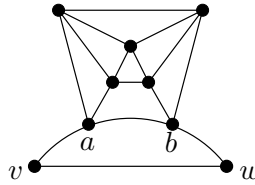
Figure 5.13: Labelled neighbourhood of a gadget subgraph in G.

thus passing through one of these vertices a second time which is a contradiction. Assuming then that ab does not lie on C_{va}, Lemma 5.12 still implies that both a and b lie on the circle. But this is impossible as neither ab nor bw are on C_{va}. Hence $C_{va} = C_{bw} := C$.

Now applying Lemma 5.12 again, we see that any circle containing both gadget and non-gadget edges must be either C or C_{ab}. We claim that the gadget can now be detached by deleting all arcs representing gadget edges. If $C = C_{ab}$, then there are no circles containing both gadget and non-gadget edges, and we are left with a circle representation of $G_2 - H_{ab}$ with a parallel edge vc. Call this resultant graph G'. If $C \neq C_{ab}$, then this deletion will cause C to be missing an arc segment. To recover a circle representation of G', all we need to do is draw this segment back in.

Detaching each of the eight segments from G_2 leaves us with a circle representation of the multigraph M, but this contradicts Theorem 5.8. $\square$

5.3 Constructions and positive results

The authors of [9] also show that 3-connectedness is a sufficient condition for a quartic planar graph to be circle representable. The proof is actually easy and relies upon the Tait graph construction from our toolkit, which they call $IL(G)$. Behind the scenes, the underlying sufficient condition for a graph G to have a circle representation is that its Tait graph TG is simple. The connection is given by the following lemma from [9].

Lemma 5.15 (Bekos and Raftopoulou [9])**.** *If G is simple and 3-connected, then TG is simple.*

Proof. Suppose there is a parallel edge in TG, say with endvertices x and y. This means that the coloured faces of G corresponding to x and y share at least two vertices, call these u and v. We claim that $\{u, v\}$ is a 2-cut in G, which would contradict 3-connectedness. To see this, fix a planar embedding of G and draw

the Tait graph so that the each edge is contained in the union of the two faces to
which it corresponds together with the common vertex. This means that u and
v are the only vertices that lie on the parallel edges between x and y, and these
parallel edges are not crossed by any edge of G. In addition, since G is simple,
there must be at least one vertex of G inside the simple closed curve formed by
the parallel edges, and also at least one vertex outside. Therefore, $\{u, v\}$ must
be a 2-cut, which can be seen in Figure 5.14(a).

Similarly, a loop in TG with endvertex z would correspond to a coloured face
of G such that some vertex, say w, appears twice in its boundary cycle. Drawing
the Tait graph as before, the loop is a simple closed curve that must have at least
one vertex of G inside and one outside since G is simple, in particular loopless.
As the curve is contained in a face of G and w is the only vertex that lies on it,
we find that w is a cutvertex as shown in Figure 5.14(b). $\square$

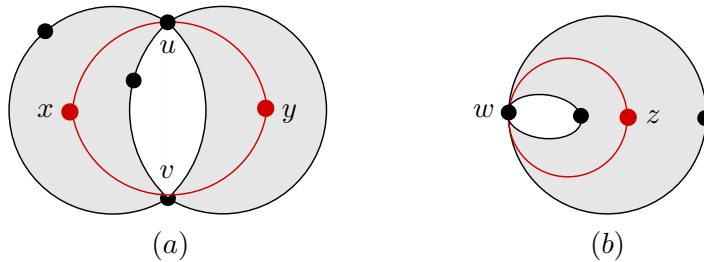

(a)   (b)

Figure 5.14: Parallel edges and loops in the TG (red) corresponds to 2-cuts and
1-cuts in G (black, only some vertices shown).

Theorem 5.16 (Bekos and Raftopoulou [9]). *Every simple 3-connected quartic
planar graph admits a circle representation.*

Proof. Let G be a graph satisfying the conditions of the theorem. Then the Tait
graph TG is simple by Lemma 5.15, and thus it has a coin representation by
Theorem 5.3. All that is left to do is to observe that a coin representation of
the Tait graph is also a circle representation of the original graph, which we have
done in Lemma 5.6. $\square$

Note that this proof produces circle representations such that all vertices are
kissing points and all circles have disjoint interiors, which seems rather wasteful
of the assumptions. A good starting point for characterising the circle repre-
sentable graphs would therefore be to come up with some constructions that
introduce these other available structures. In the special case that the Tait graph

is 3-connected, one can apply Theorem 5.4 to obtain double circle representation instead of Theorem 5.3. By definition, the double circle representation gives a (near) coin representation of the dual graph $(TG)^*$. This is not a coin representation since the circles are not all disjoint, but it is still a circle representation of $(TG)^*$ and hence of TG, since a plane graph and its dual have the same medial graph. It is therefore possible to obtain circle representations with non-disjoint circles. However, there is no way to introduce crossing points via this method. Moreover, it is easy to construct quartic planar graphs that have 1-cuts or 2-cuts and are still circle representable.

With a view to weakening the sufficient condition of 3-connectedness, we introduce a series of constructions that allow us to glue circle representations together. It is not generally obvious how to put circle representations together in a prescribed manner while preserving the underlying graphs since the arcs that we wish to make tangent may be trapped inside other faces, or we may have problematic configurations such as that in Figure 5.15.
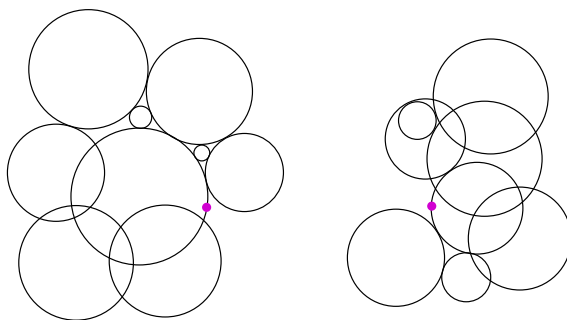

Figure 5.15: Try gluing these at the points indicated.

Our idea is to create gluing regions by transforming given circle representations using Möbius transformations, which preserve generalised circles on the extended plane, and then scaling any obstructions away or moving them off to infinity. It is most convenient to think of performing these transformations geometrically by using the inverse stereographic projection to send our representation to the sphere, rotating so that some chosen point is at the north pole, and then projecting back down again. Here is an example using 1-cuts.

Given any connected graph G that has cutvertices, there is a well-known tree-like decomposition into 2-connected components (see [42, 84]). Let the *blocks* of G be the set of its maximal 2-connected components and any bridges (together with the two endvertices). Then every edge of G is in exactly one block, and maximality ensures that any two blocks are either disjoint or intersect at a cutvertex.

If G is a quartic planar graph, then it has no bridges by Lemma 2.24 so the decomposition simplifies to having blocks be maximal 2-connected components that meet at cutvertices. We represent this as a block graph $\mathrm{Blk}(G)$, which is the intersection graph of the blocks. Explicitly, the vertices of $\mathrm{Blk}(G)$ correspond to the blocks and cutvertices of G, and there is an edge between a block vertex and a cutvertex if and only if that cutvertex is contained in the block. It is a standard result that as long as G is connected, then $\mathrm{Blk}(G)$ is a tree [84].

In order for the circle representability of blocks to make sense, a slight modification is needed to ensure that they are still quartic planar graphs. Planarity is clear. Recall that all 1-cuts in a qp-graph are balanced, so each block has a degree 2 vertex corresponding to each cutvertex that it contains. To recover a quartic planar graph, we contract one edge adjacent to each of these degree two vertices. The resulting blocks are planar and quartic but not necessarily simple. We will call these the *qp-blocks* of G. A circle representation of a *qp*-block corresponds to a circle representation of the unmodified block except that if x is a vertex of degree 2, then it is not a kissing or crossing point, but instead can be thought to lie on the arc corresponding to the edge adjacent to x that we did not contract. We will say that these arcs correspond to x.

One should view $\mathrm{Blk}(G)$ as a map telling us how to glue together circle representations of the blocks in order to obtain a circle representation of the whole graph. Figure 5.16 depicts an example where all the gluings are easily done simply by moving the individual blocks closer together. The arcs on different blocks corresponding to a cutvertex x are made tangent, so that the new kissing point represents x.
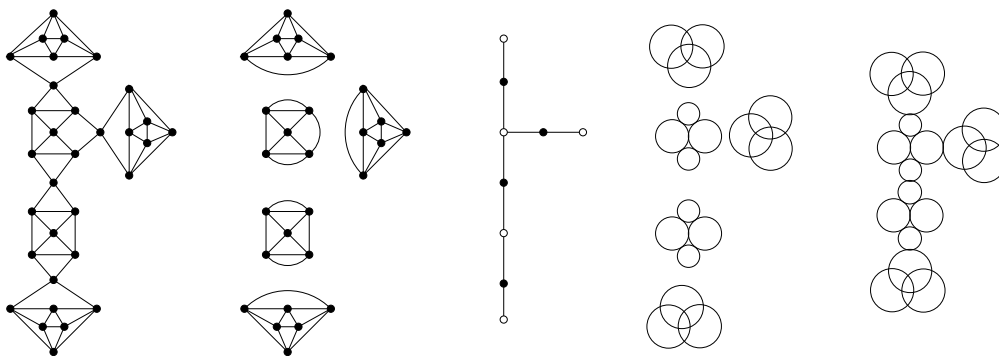

Figure 5.16: Easy block gluing. From left to right, these show G, the qp-blocks of G, $\mathrm{Blk}(G)$, circle representations of the blocks, and a circle representation of G obtained by gluing.

It turns out that this gluing is always possible, although we typically have to work a little harder to move possible obstructions out of the way.

Proposition 5.17. *Suppose that G is a quartic planar graph, and every qp-block of G is circle representable. Then G is circle representable.*

Proof. Choose any leaf of $\text{Blk}(G)$ which corresponds to some qp-block, say A, and suppose we have a path AaB in the tree where a is a cutvertex and B is another qp-block. In G, this corresponds to two maximal 2-connected subgraphs that intersect in the cutvertex a. The situation is shown in Figure 5.17(a), and the qp-blocks in (b). Then, letting A be synonymous with its circle representations, we know that there is a circle C_a in A that contains an arc corresponding to a which is drawn in magenta (c). We now apply a Möbius transformation by first using the inverse stereographic projection to view our circle representation on the sphere, then rotate so that a point on the arc is at the north pole, and project down. This means that C_a is now a straight line (circle through infinity), and all other circles are mapped to ordinary circles. Similarly, there is a circle C_a' in B that contains an arc corresponding to a, so we can repeat the procedure there.

To glue, we view the two transformed circle representations on the same plane now, and translate and rotate as required so that circles of A are disjoint to the circles of B and the lines corresponding to C_a and C_a' are parallel, as in (d). Finally, we apply another Möbius transformation sending any point inside a face (that is, not on any of the generalised circles) to infinity so that we have a collection of ordinary circles again. The images of C_a and C_a' are tangent at the point in the image of infinity. This kissing point, shown in (e), represents a so we now have a circle representation of $A \cup B$. Note that all existing tangencies and intersection points as well as the arcs between them are preserved by each transformation, and we have not created any other new vertices.

Repeating this for each path of three vertices in $\text{Blk}(G)$, we build a circle representation of the union of the qp-blocks which is precisely G. Since $\text{Blk}(G)$ is a tree in which each cutvertex has degree two, we need only glue two blocks together at a time. $\square$

In truth, the previous theorem is only a slight improvement on the Tait graph to coin representation construction in the sense that if each qp-block of G has a simple Tait graph, then one can quite easily find an embedding of the graph so that TG is actually simple as well. Unfortunately, we do not have a good description of other situations for which the hypotheses are satisfied, so this does
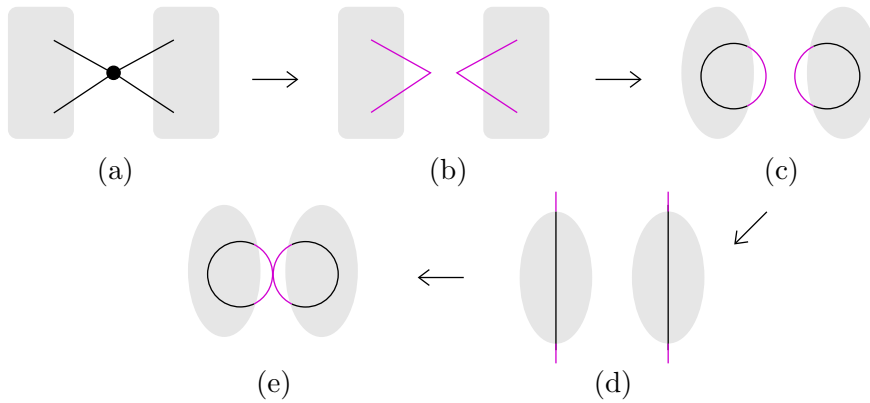
Figure 5.17: Gluing across a 1-cut.

not directly give a larger class of circle representable graphs. The advantage of our proof is that we can handle qp-blocks that have any circle representation, and in particular do not require them to use only circles with disjoint interiors or kissing points. This proposition also shows that 1-cuts are not really an obstruction to circle representability.

The natural extension would be to try to decompose 2-connected graphs into 3-connected components, and again come up with similar gluing methods. Tutte decomposition (see [84]) does this in general. However, the process involved introduces many virtual edges, and since we also need to make a modification to ensure 4-regularity, it becomes very complicated to keep track of where arcs need to be glued or identified. We do not yet have a remedy for this, although we do have some constructions that enable us to glue across 2-cuts in certain situations, so we are inclined to believe that there is some promise in these methods. We will record just one of these constructions in the case of a balanced 2-cut.

In the proof of Lemma 5.15, we saw that parallel edges in TG arise from 2-cuts in G. Upon closer inspection of that proof, only specific types of 2-cuts actually cause this issue. Suppose $\{x, y\}$ is a balanced 2-cutset in a quartic planar graph G with x adjacent to a_1, b_1, b_2 and a_2 with that cyclic ordering, and y adjacent to a_3, b_3, b_4 and a_4 also in that order, where a_i are on one side of the cut and b_i are on the other. These vertices may not all be distinct. Choosing the outer face as shown in Figure 5.18, let A_u, A_ℓ, B_u and B_ℓ denote respectively the faces incident to $a_1 x$, $a_4 y$, $b_1 x$ and $b_4 y$ that are not the outer face. If $\{A_u, A_\ell, B_u, B_\ell\}$ contains at most 3 distinct faces, then we can avoid parallel edges in TG by taking the colouring in which the outer face white, as shown Figure 5.19. The most interesting case, then, is the third one in the diagram in which it is easy
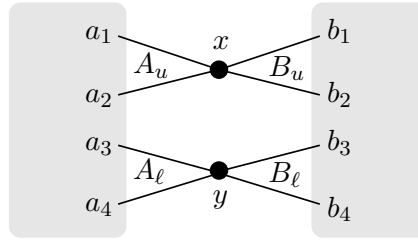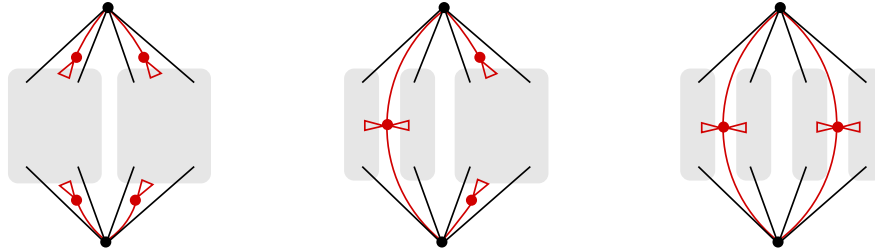
Figure 5.18: Balanced 2-cut

Figure 5.19: Faces of a balanced 2-cut, with part of TG shown in red assuming that the outer face is coloured white.

to see that both possible colourings result in parallel edges. Call the family of graphs that can be drawn in this form the *lanterns*.

Each lantern has a natural decomposition into four blocks, but as before we need to make some modification to ensure that they are 4-regular. We will do this by taking each connected component of $G - \{x, y\}$, and adding an edge between the two vertices of degree 3 in each one. It turns out that circle representations of these qp-blocks can also be glued together.

Proposition 5.18. *If a quartic planar graph G is in the family of lanterns and each of its qp-blocks is circle representable, then G is circle representable.*

Proof. By inspection, the edges of the 2-cut in a lantern look as though they should be represented by two intersecting circles. Following this intuition, our construction is illustrated in Figure 5.20, and accompanied by the following annotations.

(a) Choose a face adjacent to both vertices in the cut and draw the graph with that as the outer face.

(b) Split into qp-blocks.

(c) By assumption, these are circle representable so we now identify each block with a circle representation. Our newly created edges correspond to some

arc of a circle C_i (where $i = 1, 2, 3, 4$ corresponds to the blocks numbered from left to right in the picture) in the representation, which we are keeping track of in magenta.

(d) Viewing each block on separate planes, we apply Möbius transformations sending a point of our special arc to infinity.

(e) Taking two of our transformed blocks, say the first and third, align them so that the images of C_1 and C_3 which are lines coincide, and also all other circles representing one block are disjoint from those representing the other block. This is possible since G is finite, so each block is also finite and hence their representations consist of finitely many circles. We also do the same for the second and fourth block. It does not matter how we pair up the blocks here; the pairing we have chosen recovers the initial cyclic ordering at x and y.

(f) Again using finiteness and viewing all blocks together now, we can scale and rotate as needed so that the two lines intersect at a point p that lies between the circles representing each pair of blocks. That is, one line can be traversed so that all crossing and tangent points with circles representing C_1 occur before p whilst all those for C_3 occur after p, and similarly for C_2 and C_4 on the other line.

(g) Apply another Möbius transformation sending any point in the interior of a face to infinity, so that the images of all the generalised circles are now ordinary circles. Note that the lines map to a pair of circles that intersect at two points; the image of the p and the image of infinity.

We claim that the resulting collection of circles is precisely a circle representation of G. As before, each transformation preserves tangencies, intersections and arcs between them, and we have been careful not to create any new kissing or crossing points except the crossing point in step f which represents x, say, and the point at infinity which maps in step (g) to another crossing point that represents y.   □

There are two reasons why one should care about the above construction. Firstly, we now have a construction that introduces crossing circles. Secondly, the lanterns do not have simple Tait graphs so this a genuine improvement on the existing class of circle representable graphs. For concreteness, one can replace the condition that the qp-blocks are circle representable by the stronger condition that each (unmodified) block is 3-connected and the two vertices of degree 3 are incident to only one face in common.
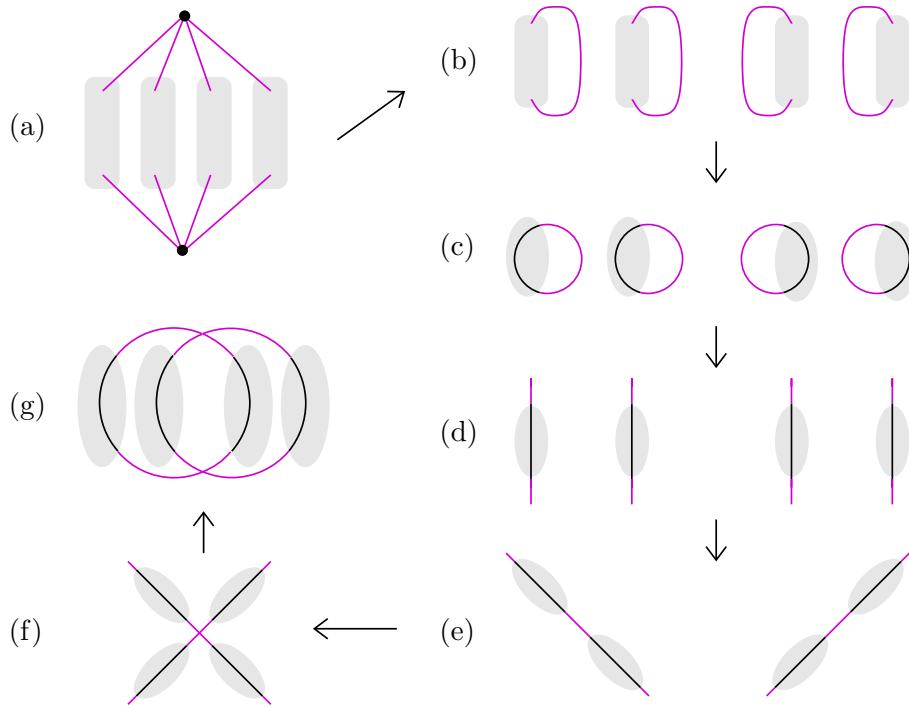
Figure 5.20: Gluing across a balanced 2-cut.

5.4 A note on kissing circle representations

If we choose to restrict our attention to kissing circle representations specifically, we are led to another interesting variation of the circle representation problem that more closely resembles the coin representation problem. Coming from the direction of the more general version, we obtain the following theorem for free.

Theorem 5.19. *All simple 3-connected quartic planar graphs admit a kissing circle representation.*

The proof is identical to that of Theorem 5.16, which we have already observed actually produces kissing circle representations.

Nonetheless, this is clearly a more restrictive type of representation and accordingly we should expect there to be fewer graphs that can be drawn in this wat. For example, it is clear that the link multigraph consisting of two vertices and four parallel edges between them does not have such a representation. Simple graphs for which we can prove this property are still not easy to find though, but the methodology we used for the general circle representation problem can also be applied successfully here. Taking the the aforementioned link as our base multigraph and adding copies of the 2-octahedral attachment, we are able to con-

struct the simple graph shown in Figure 5.21 on just 16 vertices, named the hot air balloon, which does not have a kissing circle representation.
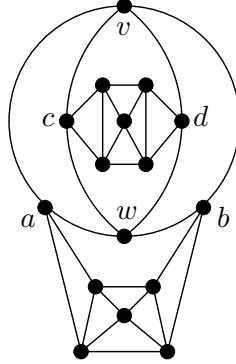

Figure 5.21: The hot air balloon does not have a kissing circle representation.

Theorem 5.20. *The hot air balloon does not have a kissing circle representation.*

Proof. This graph contains two copies of the gadget subgraph shown in Figure 5.10b. Although there are now three non-gadget vertices adjacent to the attachment vertices of each gadget rather than the two we had before, the proof of the second statement of Lemma 5.12 still goes through.

Assume that we have a circle representation of the graph, and suppose $C_{av} \neq C_{aw}$. From the preceding discussion, we may apply Lemma 5.12 to conclude that C_{av} and C_{aw} both contain vertices the a and b. However, by assumption each vertex is a kissing point. This means we have two circles that are tangent at two distinct points, which implies that $C_{av} = C_{aw}$ giving a contradiction. Hence, $C_{av} = C_{aw}$.

By symmetric arguments, we also necessarily have $C_{bw} = C_{bv}$, $C_{cv} = C_{cw}$ and $C_{dv} = C_{dw}$. Together, these four (possibly coinciding) circles contain all of the edges adjacent to v as well as all of the edges adjacent to w. Given that each vertex is a kissing point between two circles, there must be exactly two distinct circles among the elements of $\{C_{av}, C_{bv}, C_{cv}, C_{dv}\}$ that, moreover, are tangent at v and at w. This is impossible.   □

Remark 5.21. The hot air balloon counterexample also gives rise to an infinite family of graphs that do not have kissing circle representations. The base multigraph in this case consists of two vertices with four edges between them, and it is clear that this cannot be represented by kissing circles. Replacing one or both of the octahedral gadgets by any other gadget with two vertices of attachment gives a counterexample, with the proof being identical to that of Theorem 5.20.

5.5   Open problems

We have made progress toward answering the following question raised in [9], although we have not yet settled it.

Question 5.22. What is the smallest simple quartic planar graph that does not admit a circle representation?

Given the procedure we used to obtain the counterexample by going via multigraphs, we believe that the answer to the following question may shed some light on the preceding one.

Question 5.23. Can all counterexamples be derived from multigraphs that do not admit circle representations?

By 'derive' we mean that the graph can be constructed by adding some gadgets to a multigraph. Observe that if we detach the gadgets in [9], we obtain the multigraph formed from the octahedron by replacing each edge by the picture shown in 5.8b. Since that configuration could not be realised by circles, it is straightforward to deduce, using an earlier observation that digons must be realised as circle, that this multigraph does not have a circle representation either. In fact, this gives a shorter alternative proof of those counterexamples. It follows that all known counterexamples conform to the description referred to in this question.

Certainly, for any graph constructed in this way, the fact that the base multigraph is not circle representable is a good reason why the original graph should not be circle representable either. Hence, we are really asking whether there are other obstructions. An affirmative answer would reduce the problem to finding the smallest multigraph that is not circle representable, or the smallest impossible circle configuration. In that case, we would venture a guess that our configuration using graphs on 68 vertices are the smallest counterexamples as our base multigraph is already quite small and only involves 9 circles. One approach to answering this question may be to decompose the graph based on our gluing constructions, with the goal of showing that any block that is not circle-representable is a base multigraph that is not circle-representable.

It is also worth noting that we have tried the obvious alternative methods of creating a simple graph from our base multigraph using fewer attachments, and these have not produced any smaller counterexamples. An example is shown in Figure 5.22 which, despite looking deceptively similar to our counterexample, actually admits the depicted circle realisation.
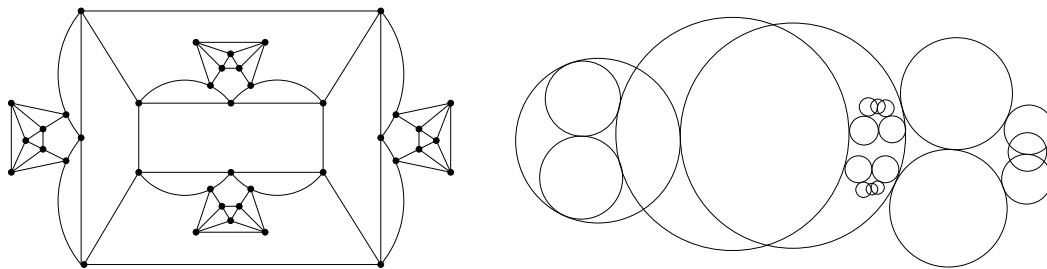
Figure 5.22: Not a counterexample.

The following bigger question also remains open.

Question 5.24. Which simple quartic graphs admit circle representations?

It is possible that an improvement to the Möbius transformation gluing constructions recorded in the previous section may yield a larger class of circle representable graphs, but we still appear to be far from a full characterisation. One could also ask the last question for any number of closely related problems, such as the restriction to kissing circle representations, or an extension to Eulerian graphs. For the latter, the natural definition of a circle representation would be to allow more than one circle to cross at a point. This actually seems to restrict the structures that are possible, so may be slightly more tractable.

Finally, our exploration of this problem was initially motivated by the observation that the 4-cycle addition operation for quartic planar graphs has a direct translation to circle representations; if G' is obtained from G by a 4-cycle addition at some vertex v and G has a circle representation, then one can obtain a circle representation of G' by adding a circle as small as necessary around the point corresponding to v. Indeed, this particular problem was mentioned in [23] accompanying a generation theorem for 3-connected quartic planar graphs, although that was while Lovász' conjecture was still open. Nonetheless, we would be interested to know if any generation theorems for other subclasses of quartic planar graphs may be applicable here.
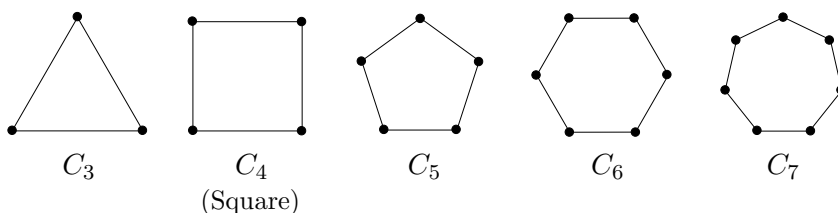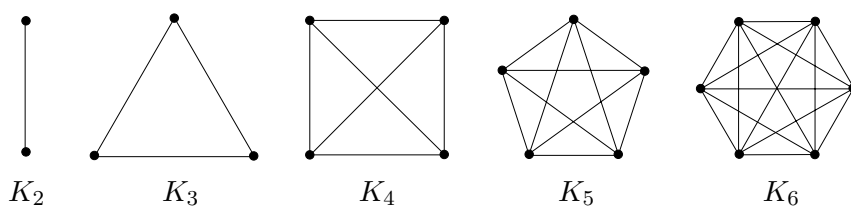
Appendix A

Gallery of graphs

The *path* P_n consist of n vertices $v_1, \ldots, v_n$ and edges $v_i v_{i+1}$ for $1 \leq i \leq n-1$.

P_2 P_3 P_4 P_5

For $n \geq 3$, adding one extra edge $v_n v_1$ to P_n gives the *n-cycle* denoted by C_n.

C_3 C_4 C_5 C_6 C_7
(Square)

The *complete graph* K_n on n vertices has an edge between every distinct pair of vertices.

K_2 K_3 K_4 K_5 K_6

The *complete bipartite graphs* $K_{m,n}$ have $m + n$ vertices, say with vertices in one colour class being $\{v_1, \ldots, v_m\}$ and the other $\{w_1, \ldots, w_n\}$. The edge set consists of $v_i w_j$ for all $1 \leq i \leq m$ and $1 \leq j \leq n$.

$K_{2,2} \cong C_4$ $K_{3,3}$ $K_{4,2}$ $K_{4,3}$
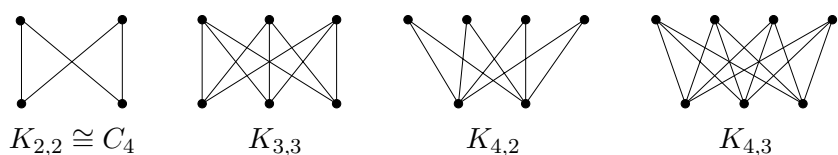
The next three triangulations are the graphs of platonic solids.

Tetrahedron $\cong K_4$ Octahedron   Icosahedron
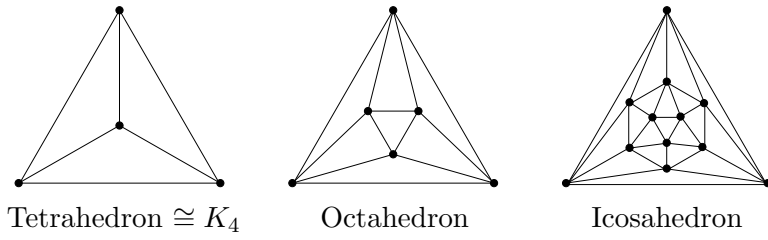
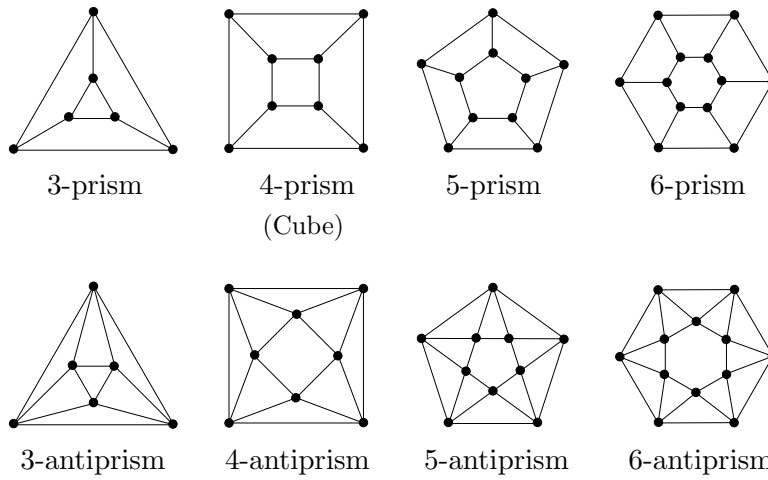Take two copies of C_n, say with vertices $(v_1, \ldots, v_n)$ and $(w_1, \ldots, w_n)$. Adding the edges $v_i w_i$ for each $1 \le i \le n$ results in the *n-prism* graph, and further adding the edges $v_i w_{i+1}$ for $1 \le i \le n$ with indices modulo n give the *n-antiprism*.

3-prism   4-prism   5-prism   6-prism
               (Cube)

3-antiprism   4-antiprism   5-antiprism   6-antiprism

The *pseudo-double wheels*, dual to the antiprisms, are quadrangulations with $n \ge 8$ vertices for n even, and consist of an $(n-3)$-cycle $(v_0, v_1, ..., v_{n-3})$, and one vertex adjacent to v_i for i even and another vertex adjacent to v_i for i odd. The smallest pseudo-double wheel is isomorphic to the cube, or 4-prism.

Cube   10-pseudo-double wheel   12-pseudo-double wheel
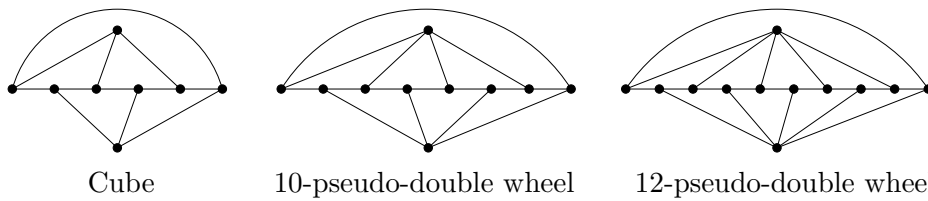
Finally, the following are all three non-isomorphic tfp-graphs on 10 vertices.

3-connected quartic planar graphs of order 10
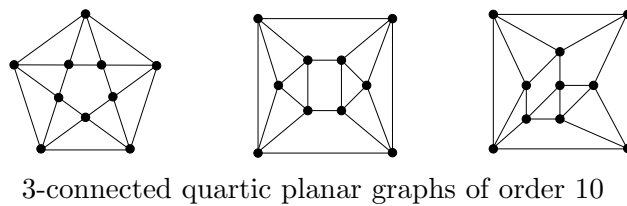
Bibliography

[1] T. Adelgren, *Triangle-free Eulerian tours in graphs with maximum degree at most 4*, Discrete Math. **138** (1995), 5–14. doi:10.1016/0012-365x(94)00183-j.

[2] M.J. Alam, D. Eppstein, M. Kaufmann, S.G. Kobourov, S. Pupyrev, A. Schulz, and T. Ueckerdt, *Contact graphs of circular arcs* (F. Dehne, J.-R. Sack, and U. Stege, eds.), Lecture Notes in Computer Science, Springer. doi:10.1007/978-3-319-21840-3_1.

[3] J. Barát and C. Thomassen, *Claw-decompositions and Tutte-orientations*, J. Graph Theory **52** (2006), 135–146. doi:10.1002/jgt.20149.

[4] D. Barnette, *On generating planar graphs*, Discrete Math. **7** (1974), 199–208. doi:10.1016/0012-365x(74)90035-1.

[5] V. Batagelj, *An inductive definition of the class of all triangulations with no vertex of degree smaller than 5*, Proc. 4th Yugoslav Seminar on Graph Theory, Novi Sad, 1983.

[6] ———, *An inductive definition of the class of 3-connected quadrangulations of the plane*, Discrete Math. **78** (1989), 45–53. doi:10.1016/0012-365x(89)90159-3.

[7] Vladimir Batagelj, *An improved inductive definition of two restricted classes of triangulations of the plane.*, Combinatorics and Graph Theory, Banach Center Publications **25** (1989), 11–18. doi:10.4064/-25-1-11-18.

[8] S. Bau, N. Matsumoto, A. Nakamoto, and L. Zheng, *Minor relations for quadrangulations on the sphere*, Graphs and Combin. **31** (2014), 2029–2036. doi:10.1007/s00373-014-1489-y.

[9] M. A. Bekos and C. N. Raftopoulou, *On a conjecture of Lovász on circle-representations of simple 4-regular planar graphs*, J. Comput. Geom. **6** (2015), no. 1, 138–149. doi:10.20382/jocg.v6i1a1.

[10] J. Bensmail, A. Harutyunyan, T-N. Le, and S. Thomassé, *Edge-partitioning a graph into paths: beyond the Barat-Thomassen conjecture*, Combinatorica (2016). arXiv:1507.08208.

[11] J.A. Bondy, *Perfect path double covers of graphs*, J. Graph Theory **14** (1990), 259–272. doi:10.1002/jgt.3190140213.

[12] F. Bories, J.-L. Jolivet, and J.-L. Fouquet, *Construction of 4-regular graphs*, Ann. Discrete Math. **17** (1983), 99–118. (In French) doi:10.1016/s0304-0208(08)73378-4.

[13] R. Bowen and S. Fisk, *Generation of triangulations of the sphere*, Math. Comp. **21** (1967), 250–252. doi:10.2307/2004172.

[14] G.R. Brightwell and E. Scheinerman, *Representations of planar graphs*, SIAM J. Disc. Math. **6** (1993), 214–229.

[15] G. Brinkmann and M. Deza, *Lists of face-regular polyhedra*, J. Chem. Inform. Comput. Sci. **40** (2000), 530–541. doi:10.1021/ci990034u.

[16] G. Brinkmann, J. Goedgebeur, and B.D. McKay, *The generation of fullerenes*, J. Chem. Inf. Model. **52** (2012), 2910–2918. doi:10.1021/ci3003107.

[17] G. Brinkmann, S. Greenberg, C. Greenhill, B.D. McKay, R. Thomas, and P. Wollan, *Generation of simple quadrangulations of the sphere*, Discrete Math. **305** (2005), 33–54. doi:10.1016/j.disc.2005.10.005.

[18] G. Brinkmann, T. Harmuth, and O. Heidemeier, *The construction of cubic and quartic planar maps with prescribed face degrees*, Discrete Appl. Math. **128** (2003), no. 2-3, 541–554. doi:10.1016/s0166-218x(02)00549-8.

[19] G. Brinkmann, C. Larson, J. Souffriau, and N. Van Cleemput, *Construction of planar 4-connected triangulations*, Ars Math. Contemp. **9** (2015), no. 2, 145–149. doi:10.26493/1855-3974.459.c4d.

[20] G. Brinkmann and B.D. McKay, `plantri` (software). `http://users.cecs.anu.edu.au/~bdm/plantri`.

[21] _____, *Construction of planar triangulations with minimum degree* 5, Discrete Math. **301** (2005), no. 2-3, 147–163. doi:10.1016/j.disc.2005.06.019.

[22] _____, *Fast generation of planar graphs*, MATCH Commun. Math. Comput. Chem. **58** (2007), 323–357.

[23] H.J. Broersma, A.J.W. Duijvestijn, and F. Göbel, *Generating all 3-connected and 4-regular planar graphs from the octahedron graph*, J. Graph Theory **17** (1993), no. 5, 613–620. doi:10.1002/jgt.3190170508.

[24] J.W. Butler, *A generation procedure for the simple* 3-*polytopes with cyclically* 5-*connected graphs*, Canad. J. Math. **26** (1974), 686–708. doi:10.4153/cjm-1974-065-6.

[25] H. de Fraysseix, P.O. de Mendez, and P. Rosenstiehl, *On triangle contact graphs*, Combinatorics, Probability and Computing **3** (1994), no. 2, 233–246. doi:10.1017/s0963548300001139.

[26] R. Diestel, *Graph Theory*, 3rd ed., GTM, Springer-Verlag, Berlin Heidelberg, 2005.

[27] G. Ding and J. Kanno, *Splitter theorems for 4-regular graphs*, Graphs and Combin. **26** (2010), no. 3, 329–344. doi:10.1007/s00373-010-0916-y.

[28] G. Ding, J. Kanno, and J. Su, *Generating 5-regular planar graphs*, J. Graph Theory **61** (2009), no. 3, 219–240. doi:10.1002/jgt.20377.

[29] C.A. Duncan, E.R. Gansner, Y.F. Hu, M. Kaufmann, and S.G. Kobourov, *Optimal polygonal representation of planar graphs*, Algorithmica **63** (2012), no. 3, 672–691. arXiv:1104.1482.

[30] V. Eberhard, *Zur morphologie der polyeder*, Teubner, 1891.

[31] P. Erdös, A. Rényi, and V.T. Sós, *Combinatorial theory and its applications*, North-Holland, Amsterdam, 1970.

[32] B. Faller and B.D. McKay, private communication.

[33] I. Fáry, *On straight line representation of planar graphs*, Acta Univ. Szeged. Sect. Sci. Math. **11** (1948), 229–233.

[34] Stefan Felsner, *Rectangle and Square Representations of Planar Graphs*, Thirty Essays on Geometric Graph Theory, Springer. doi:10.1007/978-1-4614-0110-0_12.

[35] O.D. Friedrichs, A.W.M. Dress, A. Müller, and M.T. Pope, *Polyoxometalates: a class of compounds with remarkable topology*, Mol. Eng., 9–28. doi:10.1007/978-94-011-0920-8_2.

[36] E. Fuchs and L. Gellert, *Reducing quadrangulations of the sphere and the projective plane*. arXiv:1606.07662.

[37] C. Godsil and G. Royle, *Algebraic Graph Theory*, Springer-Verlag, New York, 2001.

[38] D. Goncalves, B. Léveque, and A. Pinlou, *Triangle contact representations and duality*, Discrete Comput. Geom. **48** (2012), no. 1, 239–254. doi:10.1007/s00454-012-9400-1.

[39] J.L. Gross and T.W. Tucker, *Topics in Topological Graph Theory*, Cambridge University Press, New York, 2009.

[40] B. Grünbaum, *Convex Polytopes*, Vol. 2nd, Spinger-Verlag, New York, 2003.

[41] R. Häggkvist, *Decompositions of regular bipartite graphs* (J Siemons, ed.), Surveys in Combinatorics, LMS Notices 141, Cambridge University Press, 1989. doi:10.1017/cbo9781107359949.007.

[42] F. Harary, *Graph theory*, Addison-Wesley Publishing Company, 1969.

[43] M. Hasheminezhad, B.D. McKay, and T. Reeves, *Recursive generation of simple planar 5-regular graphs and pentangulations*, J. Graph Algorithms Appl. **15** (2011), no. 3, 417–436. doi:10.7155/jgaa.00232.

[44] Z.-X. He and O. Schramm, *Fixed points, Koebe uniformization and circle packings*, Annals of Mathematics **137** (1993), no. 2, 369–406. doi:10.2307/2946541.

[45] P. Hlinený, *Contact graphs of curves*, 3rd Int. Symp. Graph Drawing (GD1995), Springer, Lect. Notes Comput. Sc. (1996). doi:10.1007/bfb0021814.

[46] S. Jimbo, *On the Eulerian recurrent lengths of complete bipartite graphs and complete graphs*, IOP Conf. Ser.: Mater. Sci. Eng **58** (2014). doi:10.1088/1757-899x/58/1/012019.

[47] ———, *The NP-completeness of Eulerian recurrent length for 4-regular Eulerian graphs*, 4th International Conference on Artificial Intelligence with Applications in Engineering and Technology (2014). doi:10.1109/icaiet.2014.34.

[48] S. Jimbo, Y. Oshie, and K. Hashiguchi, *The NP-completeness of Eulerian Recurrent Length*, Kyoto University Departmental Bulletin (2005).

[49] M. Jooyandeh, *Recursive algorithms for generation of planar graphs*, Ph.D Thesis, Australian National University, 2014.

[50] J. Liu K. Heinrich and M. Yu, *P₄-decompositions of regular graphs*, J. Graph Theory **31** (1999), no. 2, 135–143.

[51] J. Kanno, N. Matsumoto, J. Su, and K. Yamamoto, *Diagonal transformations in pentangulations on the sphere*, Ars Combin. **135** (2017), no. 299–310.

[52] R. Kápolnai, G. Domokos, and T. Szabó, *Generating spherical multiquadrangulations by restricted vertex splittings and the reducibility of equilibrium classes*, Period. Polytech., Electr. Eng. **56** (2012), no. 1, 11–20. doi:10.3311/ppee.7074. arXiv:1206.1698.

[53] P. Koebe, *Kontaktprobleme der konformen abbildung*, Berichte über die Verhandlungen d. Sächs. Akad. d. Wiss., Math.–Phys. Klasse **88** (1936), 141–164.

[54] A. Kotzig, *From the theory of finite regular graphs of degree three and four*, Casopis Pest Mat. **82** (1957), 76–92.

[55] T.-N. Le, *Locally self-avoiding Eulerian tours*, J. Combin. Theory Ser. B (2018). doi:10.1016/j.jctb.2018.08.008. arXiv:1611.07486.

[56] J. Lehel, *Generating all 4-regular planar graphs from the graph of the octahedron*, J. Graph Theory **5** (1981), 423–426. doi:10.1002/jgt.3190050412.

[57] C.E. Leiserson, *Area-efficient graph layouts (for VLSI)*, Proc. 21st Annual IEEE Symposium on Foundations of Computer Science, 1980, pp. 270–281. doi:10.1109/sfcs.1980.13.

[58] P. Manca, *Generating all planar graphs regular of degree four*, J. Graph Theory **3** (1979), 357–464. doi:10.1002/jgt.3190030406.

[59] B.D. McKay, *Isomorph-free exhaustive generation*, J. Algorithms **26** (1998), no. 2, 306–324. doi:10.1006/jagm.1997.0898.

[60] G.L. Miller and W. Thurston, *Separators in two and three dimensions*, Proc. 22th ACM STOC (1990), 300–309. doi:10.1145/100216.100255.

[61] B. Mohar and C. Thomassen, *Graphs on Surfaces*, Johns Hopkins University Press, 2001.

[62] M. Nagashima, A. Nakamoto, S. Negami, and Y. Suzuki, *Generating 3-connected quadrangulations on surfaces*, Ars Combin. **116** (2014), 371–384.

[63] A. Nakamoto, *Irreducible quadrangulations of the Klein bottle*, Yokohama Math. J. **43** (1990), 125–139.

[64] _____, *Diagonal transformations in quadrangulations of surfaces*, J. Graph Theory **21** (1996), 289–299. doi:10.1002/(sici)1097-0118(199603)21:3¡289::aid-jgt3¿3.0.co;2-m.

[65] _____, *Irreducible quadrangulations of the torus*, J. Combin. Theory Ser. B **67** (1996), 183–201. doi:10.1006/jctb.1996.0040.

[66] A. Nakamoto and S. Negami, *Diagonal transformations of graphs on closed surfaces*, Sci. Rep. Yokohama Nat. Univ. Sect. I Math. Phys. Chem. **40** (1994), 71–97.

[67] A. Nakamoto and K. Ota, *Note on irreducible triangulations of surfaces*, J. Graph Theory **20** (1995), no. 2, 227–233. doi:10.1002/jgt.3190200211.

[68] N. Oksimets, *Euler Tours of Maximum Girth in K_{2n+1} and $K_{2n,2n}$*, Graphs and Combin. **21** (2005), 107–118. doi:10.1007/s00373-004-0578-8.

[69] Ø. Ore, *The Four-Color Problem*, Academic Press, 1967.

[70] H. Tang P. A. Pevzner and M.S. Waterman, *An Eulerian path approach to DNA fragment assembly*, Proc. National Academy of Sciences **98** (2001), no. 17, 9748–9753. doi:10.1073/pnas.171285098.

[71] M.S. Rahman, T. Nishizeki, and S. Ghosh, *Rectangular drawings of planar graphs*, J. Algorithms **50** (2004), no. 1, 62–78. doi:10.1016/s0196-6774(03)00126-3.

[72] J.L. Ramírez-Alfonsín, *The spread of K_n*, Discrete Math. **175** (1997), no. 1, 221–229. doi:10.1016/s0012-365x(96)00152-5.

[73] N. Robertson and P.D. Seymour, *Graph Minors. XX. Wagner's conjecture*, J. Combin. Theory Ser. B **92** (2004), no. 2, 325–357. doi:10.1016/j.jctb.2004.08.001.

[74] H. Sachs, *Coin graphs, polyhedra, and conformal mapping*, Discrete Math. **134** (1994), 133–138. doi:10.1016/0012-365x(93)e0068-f.

[75] H. Sachs, M. Stiebitz, and R.J. Wilson, *An historical note: Euler's Königsberg letters*, J. Graph Theory **12** (1988), no. 1, 133–139. doi:10.1002/jgt.3190120114.

[76] O. Schramm, *Square tilings with prescribed combinatorics*, Israel Jour. Math. **84** (1993), 97–118. doi:10.1007/bf02761693.

[77] E. Steinitz and H. Rademacher, *Vorlesungen über die Theorie der Polyeder*, Springer, Berlin, 1934. (In German).

[78] T. Sulanke, *Generating irreducible triangulations of surfaces.* arXiv:math/0606687.

[79] _____, *Source for surftri and lists of irreducible triangulations.* Available at `http://hep.physics.indiana.edu/~tsulanke/graphs/surftri/`.

[80] _____, *Note on the irreducible triangulations of the Klein bottle*, J. Combin. Theory Ser. B **96** (2006), no. 6, 964–972. arXiv:math/0407008.

[81] _____, *Generating maps on surfaces*, Discrete Comput. Geom. **57** (2017), no. 2, 335–356. arXiv:1509.06412.

[82] Y. Suzuki, *Cube-contractions in 3-connected quadrangulations*, Ars Math. Contemp. **10** (2016), 281–290. doi:10.26493/1855-3974.552.bf3.

[83] W.T. Tutte, *A theory of 3-connected graphs*, Nederl. Akad. Wetensch. Proc. Ser. A **64** (1961), 441–455.

[84] _____, *Graph theory*, Cambridge University Press, Cambridge, 2001.

[85] W.T. Tutte and C.A.B. Smith, *On unicursal paths in a network of degree 4*, Amer. Math. Monthly **48** (1941), no. 4, 233–237. doi:10.2307/2302716.

[86] K. Wagner, *Bemekungen zum Vierfarbenproblem*, J. der Deut. Math. **46** (1936), no. 1, 26–32.