Università degli Studi di Padova
Dipartimento di Matematica "Tullio Levi-Civita"

Corso di Laurea Magistrale in
MATEMATICA

Tesi di Laurea Magistrale

Coperture di grafi attraverso cicli

Graph coverings with cycles

Relatore:
Prof. Marco Di Summa

Laureando: Paolo Franchi
Matricola: 1156472

28 Settembre 2018

Contents

1 Introduction **5**

1.1 Definitions . . . 6

1.1.1 Graphs, subgraphs . . . 6

1.1.2 Walks, paths, tours, cycles . . . 7

1.1.3 Connectivity, trees . . . 7

1.1.4 Matchings, cuts . . . 8

1.2 Infinite graphs . . . 9

1.3 Preliminaries . . . 10

1.3.1 Decompositions of walks, even graphs . . . 10

1.3.2 Paths in trees . . . 13

1.3.3 Finite and countable versions of Nash-Williams' decomposition theorem . . . 16

2 Decompositions into cycles **17**

2.1 Cycles covering each edge at most countably many times . . . 17

2.1.1 Countable covering cycle collections . . . 17

2.1.2 The first construction . . . 19

2.1.3 The cycle covering theorem . . . 24

2.2 Cut-faithful countable subgraphs . . . 25

2.3 Nash-Williams' cycle-decomposition theorem . . . 30

2.4 Summary of the important results . . . 31

3 Coverings with cycles **33**

3.1 Reduction to snarks . . . 33

3.2 Cycles covering each edge at most finitely many times . . . 35

3.2.1 Finite covering cycle collections . . . 35

3.2.2 The second construction . . . 36

3.2.3 The covering result . . . 39

3.2.4 Directed graphs . . . 39

3.3 Graphs with a cycle double cover . . . 40

3.3.1 Every 4-edge-connected graph has a cycle double cover . . 40

3.3.2 Every Hamiltonian graph has a cycle double cover . . . 41

3.3.3 1-factorable graphs . . . 41

3.3.4 Final remarks . . . 42

Chapter 1

Introduction

A decomposition of a graph is a collection of pairwise edge-disjoint subgraphs such that each edge of the graph is contained in exactly one of the subgraphs. We are interested in the problem of decomposing a graph into cycles. A finite graph has such a decomposition if and only if it is Eulerian, i.e. if and only if each vertex has even degree. The result can be extended to infinite graphs and is known as Nash-Williams' cycle-decomposition theorem. The condition on the degrees has to be changed to a different one (we ask that all finite cuts are even); they are equivalent to each other in the finite case, see Section 1.3. In Chapter 2, we give a proof of this theorem.

If the graph does not have a decomposition into cycles, we study a similar problem: can the graph be covered by cycles? A covering is again a collection of subgraphs but now an edge may be in more than one of the subgraphs. It has been conjectured that every 2-edge-connected graph can be covered by cycles so that each edge is in exactly two cycles. This result, known as the cycle double cover conjecture, remains an open problem in graph theory and is discussed in Chapter 3.

In Section 3.2 we discuss the conjecture and prove a weaker result: we show that every 2-edge-connected graph has a collection of cycles such that each edge is in at least one and at most 7 cycles in the collection.

Part of this dissertation is based on the study of the article "Nash-Williams' cycle-decomposition theorem" by C. Thomassen [13]. The results and the proofs presented in Chapter 2 and Section 3.2 were given in the referenced article.

Some of the original work done for this thesis, with regard to Chapter 2 and Section 3.2, includes the expansion of some proofs, a refinement of the organization of contents and some readability improvements.

In Section 3.3, we give an elementary proof of the fact that every 4-edge-connected graph has a cycle double cover. We also show that Hamiltonian graphs and k-regular, 1-factorable, 2-edge-connected graphs have a cycle double cover.

Such proof was written by the author of the present thesis; although it seems to be hard to completely determine its originality, we could not find any published reference explicitly mentioning this method and applying it to 4-edge-connected graphs, to Hamiltonian graphs and to 1-factorable graphs.

In Section 3.1 we give a brief overview of some properties of a potential minimal counterexample to the cycle double cover conjecture, presenting results by Jaeger [6] and Huck [4].

1.1 Definitions

1.1.1 Graphs, subgraphs

A *graph* G is an ordered pair $G = (V, E)$ where V is any set and E is a set of 2-element subsets of V. The elements of V are called *vertices* of G, the elements of E *edges* of G.

If G is a graph, we denote by $V(G), E(G)$ its vertex and edge sets, respectively. Given two vertices $v, w \in V(G)$, we say that they are *neighbors* if $\{v, w\} \in E(G)$; we also say that v, w are the *endpoints* of the edge $\{v, w\}$.

A graph is *finite* if $V(G) \cup E(G)$ is a finite set; a graph is *infinite* if it is not finite. In other words, an infinite graph may have infinitely many vertices, infinitely many edges or both.

If v is a vertex in a graph G, the *degree* of v (in G), denoted by $d_G(v)$, or $d(v)$ if there is no ambiguity, is the cardinality (possibly infinite) of the set of neighbors of v.

A graph is said to be *k-regular* if every vertex has degree k. A graph is *cubic* if it is 3-regular. A finite graph is *even* if every vertex has even degree.

Let G, G' be graphs. We say that G' is a *subgraph* of G, and write $G' \subseteq G$, if $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$.

If $G' \subseteq G$ and G' contains all the edges of G with both endpoints in $V' = V(G')$, then G' is an *induced subgraph* of G; we write $G' = G[V']$. Thus if $S \subseteq V$ is any set of vertices, $G[S]$ denotes the graph on S whose edges are precisely the edges of G with both endpoints in S.

If A is any set of edges of G, we denote by $G - A$ the subgraph of G consisting of all edges of G except those in A. We refer to this operation as the *removal of the edges* in A from the graph G.

If U is any set of vertices of G, we write $G - U$ for $G[V \setminus U]$. In other words, $G - U$ is obtained from G by deleting all the vertices in U and the edges incident with them.

Two graphs G, G', defined on the same vertex set, are said to be *edge-disjoint* if $E(G) \cap E(G') = \varnothing$.

A *decomposition* of a graph is a collection of pairwise edge-disjoint subgraphs such that each edge of the graph is contained in exactly one of the subgraphs.

A *covering* of a graph is a collection of subgraphs such that each edge of the graph is contained in at least one of the subgraphs. In the case of coverings, the subgraphs are not necessarily pairwise edge-disjoint.

Given a set V, a *multigraph* G is an ordered pair $G = (E, g)$ where E is a set and g is a function:

$$g : E \to \{\{v, w\} : v, w \in V, v \neq w\}.$$

E is the set of (multiple) edges of G; if $e \in E$ and $g(e) = \{v, w\}$, we say that e is an edge between v and w. As before, v, w are the endpoints of e; given two vertices, the edge that has those vertices as endpoints may not be determined uniquely. If g is injective, G is a graph in the sense defined above.

We denote a multiple edge consisting of n edges joining the same two vertices by *n-edge* (or *single*, *double*, *triple edge* if $n = 1, 2, 3$ respectively).

All the previous notions can be extended to multigraphs. We will not distinguish between graphs and multigraphs, unless explicitely mentioned, and use the word graph for both.

1.1.2 Walks, paths, tours, cycles

A *walk* in G is a finite sequence $W = v_0e_1v_1e_2v_2\dots e_kv_k$ whose terms are alternately vertices and edges, such that the edges $e_1, e_2, \dots, e_k$ are distinct and for $i \in \{1, 2, \dots, k\}$, the ends of e_i are v_{i-1} and v_i. We say that W is a walk *from* v_0 *to* v_k, and call these two vertices the *endpoints* of W. The other vertices, $v_1, v_2, \dots, v_{k-1}$ are the *intermediate vertices* of W. The number k is the *length* of W.

If the vertices $v_0, v_1, \dots, v_k$ are distinct, W is called a *path*.

A *closed walk*, or *tour*, is a walk with positive length and such that $v_0 = v_k$. A *cycle* is a closed walk in which there is no repetition of vertices (other than $v_0 = v_k$). In the case of multigraphs, a double edge is considered to be a cycle of length 2.

We will refer to walks, paths, tours and cycles by the sequence of their vertices.

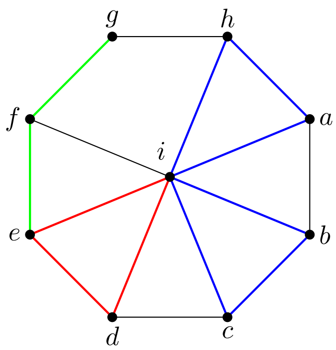

Figure 1.1: efg (green) is a path, $aibciha$ (blue) is a tour, $deid$ (red) is a cycle.

The *distance* between two vertices in a graph is the length of the shortest path connecting them. If no such path exists, the distance is conventionally defined as infinite.

1.1.3 Connectivity, trees

A graph G is *connected* if for any two vertices v, w of G, there is a path in G from v to w. A graph is *disconnected* if it is not connected. In a connected (possibly infinite) graph, the distance is finite for all pairs of vertices.

A maximal connected subgraph of G is called a *component* of G.

A graph is k-*connected* if it contains at least $k+1$ vertices and after deleting any $k-1$ vertices, the resulting subgraph is connected.

A graph is k-*edge-connected* if after deleting any $k-1$ edges, the resulting subgraph is connected. In a k-edge-connected graph, every cut contains at least k edges.

A *bridge* is an edge whose removal disconnects the graph. A connected graph is 2-edge-connected if and only if it does not have any bridge.

A *block* (or *biconnected component*) is a maximal 2-connected subgraph. Different blocks meet at *cut vertices*, vertices whose removal disconnects the graph. Bridges and isolated vertices are blocks of the graph.

A connected graph with no odd cut cannot have any bridge and is therefore 2-edge-connected. The converse is false: the graph in Figure 1.2 is 2-edge-connected but has a cut of size 3.

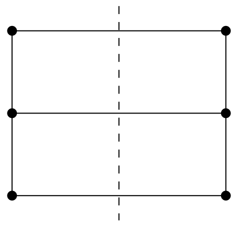

Figure 1.2: A 2-edge-connected graph with an odd cut.

A *tree* is a connected graph containing no cycle.

The following statements are all equivalent for a graph G:

1. G is a tree.
2. (Path uniqueness) For any two vertices $v, w \in V$, there exists a unique path from x to y.
3. (Minimal connected graph) G is connected and deleting any of its edges results in a disconnected graph.
4. (Maximal graph without cycles) G contains no cycles, and any graph obtained from G by adding an edge contains a cycle.

Moreover, if G is finite, then G is a tree if and only if:

5. (Euler's formula) G is connected and $|V(G)| = |E(G)| + 1$.

A *forest* is a graph whose components are trees.

If we remove a maximal collection of pairwise edge-disjoint cycles from a graph, what remains is a forest.

Let $G = (V, E)$ be a graph. A tree of the form $T = (V, E')$, where $E' \subseteq E$, is called a *spanning tree* of G. In other words, a spanning tree is a subgraph of G that is a tree and contains all vertices of G.

Every finite graph has a spanning tree. See Section 1.2 for some results on spanning trees in general graphs.

1.1.4 Matchings, cuts

A *matching* is a set of edges in which no two share any vertex.

Let G be a graph and $H \subseteq G$ a subgraph; the *boundary* of H is the set of vertices in H having a neighbor outside H.

If we add a collection of edges to G such that the added edges form a matching, then we call that set of edges an *external matching*.

A matching is said to be *perfect* if it contains every vertex of the graph.

A *cut* is a partition of the vertices of a graph into two sets, which are called the *sides* of the cut. A *cut-set* is the set of all edges having one endpoint in each side of a cut. We will often refer to cut-sets simply as cuts.

A *k-cut* is a cut consisting of k edges. A cut is said to be *trivial* if one of its sides consists of only one vertex of the graph.

A cut is *minimal* if it contains no other cut as a proper subset.

In a connected graph G, a cut is minimal if and only if its removal produces exactly two connected components, that is if $G[A], G[B]$ are both connected (where A, B are the sides of the cut).

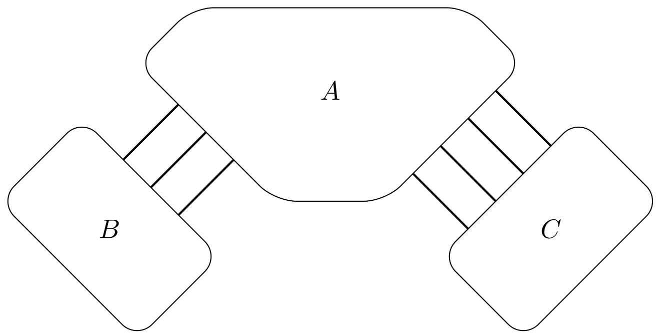

Figure 1.3: The cut between A and $B \cup C$ is not minimal.

A subgraph H of a graph G is said to be *cut-faithful* if every finite minimal cut in H is a cut in G. A cut-faithful subgraph preserves the edge-connectivity; see the proof of Theorem 13.

An *edge-coloring* of a graph is an assignment of colors to the edges of the graph such that no two adjacent edges have the same color. A graph is said to be *k-edge-colorable* if it can be edge-colored with k colors. The *chromatic index* of a graph, denoted by $\chi'(G)$, is the minimum number of colors required to edge-color the graph.

1.2 Infinite graphs

A graph G is *countable* if $V(G) \cup E(G)$ is a finite or countably infinite set.

A graph is called *locally finite* [resp. *locally countable*] if every vertex has finite [resp. countable] degree.

Theorem 1. *Every connected locally countable infinite graph is a countable graph.*

Proof. Let G be a connected locally countable infinite graph, v one of its vertices, $A_0 = \{v\}$, A_1 the set of vertices adjacent to v, A_2 the set of all vertices adjacent to a vertex of A_1, and so on. By assumption A_1 is countable. Now A_2 is a countable union of countable sets, hence it is also countable; then also A_3, A_4 and so on. The union $\bigcup_{i=0}^{\infty} A_i$ is countable and contains every vertex of G (because it is connected). G can only contain countably many vertices. Finally, since it is locally countable, G can only contain countably many edges. $\square$

In particular, every component of a locally countable infinite graph is a countable graph.

A countable graph has only countably many finite edge-sets and hence countably many finite minimal cuts.

Remarks on the Axiom of Choice

In the theory of infinite graphs, some results depend on the Axiom of Choice; here are some of its equivalent formulations:

1. Axiom of Choice: For every set X of non-empty sets there exists a choice function on X;

2. Zorn's lemma: Every non-empty partially ordered set in which every chain has an upper bound has a maximal element;

3. Well-ordering theorem: Every set can be well-ordered;

4. Every connected graph has a spanning tree (see [10] for a proof of the equivalence).

The well-ordering theorem appears explicitly in the proof of Theorem 5. In Section 2.1 we will take a spanning tree of a general graph.

Zorn's lemma is also used several times in order to obtain maximal elements, typically maximal collections of paths. The general idea of its use is as follows.

We start with a set of vertices in a graph and consider all sets S of pairwise edge-disjoint paths having endpoints in these vertices and possibly satisfying some additional property (e.g. having length 2). The set X of all these sets S is partially ordered by inclusion. Moreover X has the property required by Zorn's lemma: in fact, given a chain (i.e. a totally ordered subset) in X, the union of the chain is an upper bound to it. Therefore X has a maximal element by Zorn's lemma.

Some examples of this are in the proof of Theorem 5 and in the construction shown in Subsection 2.1.2.

1.3 Preliminaries

1.3.1 Decompositions of walks, even graphs

Walks and tours (as opposed to paths and cycles) may have repetitions of vertices; such repetitions create cycles, which can be removed from the (closed) walk as follows:

- A tour can be decomposed into pairwise edge-disjoint cycles.

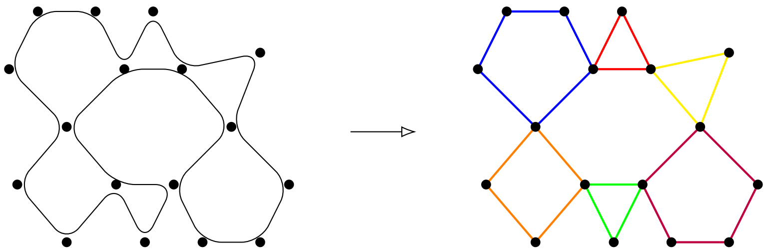

Figure 1.4: A tour decomposed into cycles.

- A walk w can be decomposed into cycles and a path with the same ends as w.

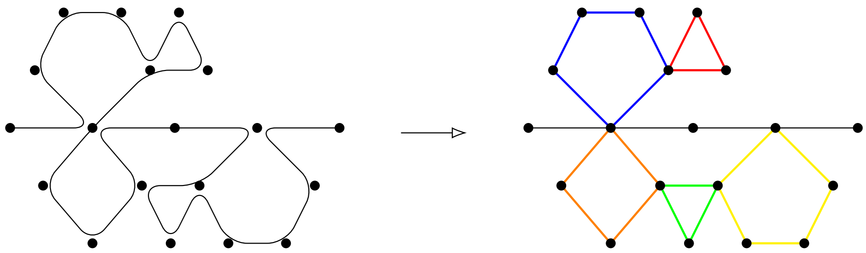

Figure 1.5: A walk decomposed into cycles and a path with the same ends.

It is well known that a finite graph is Eulerian if and only if every vertex has even degree. Given an Eulerian tour of a graph, we can decompose it into pairwise edge-disjoint cycles. Conversely, if a graph is decomposed into cycles, then every vertex has necessarily even degree (in fact the degree of a vertex v equals twice the number of cycles in which v is contained).

We want to show that, for finite graphs, there is an additional condition which is equivalent to being Eulerian.

Theorem 2. *In a finite graph, every vertex has even degree if and only if every cut is even.*

Proof. Given a vertex v, the size of the cut with sides $\{v\}$ and $V \setminus \{v\}$ is even by assumption and equals the degree of v.

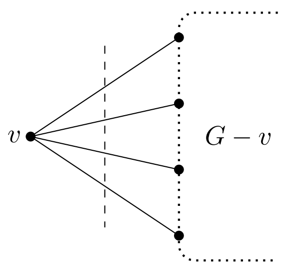

Figure 1.6: The degree of v equals the size of the cut between $\{v\}$ and $V \setminus \{v\}$.

Conversely, suppose that every vertex has even degree and consider a cut with sides A, B. The numbers $\sum_{v\in A} d(v), \sum_{v\in B} d(v)$ are both even because sums of even numbers.

Let n_A [resp. n_B] be the number of edges with both endpoints in A [resp. in B]. Also let n_C be the number of edges in the cut. Then

$$\begin{aligned} n_C &= \sum_{v\in A} d(v) - 2n_A \\ &= \sum_{v\in B} d(v) - 2n_B. \end{aligned}$$

In both cases, the expression on the right, hence also n_C, is an even number. □

Therefore, if G is a finite graph, the following statements are all equivalent:

1. G is even;
2. G is Eulerian;
3. G has a decomposition into cycles;
4. Every cut in G is even.

For infinite graphs, the above conditions may not be equivalent or even make sense:

1. It is possible to check whether a graph is even only if it is locally finite; since we will deal with graphs that are generally not locally finite, we are not going to discuss this condition.

2. Recall that tours, including Eulerian tours, are finite sequence of edges; therefore no graph with infinitely many edges can have an Eulerian tour. Notice that there are easy examples of graphs having all vertices of even degree but with no Eulerian tour, such as the infinite trail pictured in Figure 1.7.

Figure 1.7: In the infinite trail, every vertex has degree 2.

3. It makes sense to ask whether a general graph has a decomposition into cycles. This condition appears in Nash-Williams' cycle decomposition theorem (Theorem 14).

4. The cuts in a graph may in general be infinite, hence the condition "Every cut is even" makes sense only if there is at least one finite cut. We need to change it to "Every finite cut is even", or equivalently, "There is no finite odd cut". Then it becomes equivalent to having a decomposition into cycles.

If G_1, G_2 are subgraphs of a graph G, their *symmetric difference*, denoted by $G_1 \triangle G_2$, is the subgraph consisting of all edges of G contained in one of G_1, G_2 but not the other.

In the next Theorem, we show that the symmetric difference of two even graphs is also an even graph.

Theorem 3. *Let G be a graph, G_1, G_2 subgraphs of G. Suppose that G_1, G_2 are even and their union is G. Then $G_1 \triangle G_2$ is an even graph.*

Proof. Let v be a vertex of G. We want to show that $d_{G_1 \triangle G_2}(v)$ is even. Let n be the number of edges having v as an endpoint and belonging to both G_1 and G_2. Then

$$d_{G_1 \triangle G_2}(v) = (d_{G_1}(v) - n) + (d_{G_2}(v) - n) = d_{G_1}(v) + d_{G_2}(v) - 2n.$$

By assumption, the numbers $d_{G_1}(v), d_{G_2}(v)$ are even, hence also $d_{G_1 \triangle G_2}(v)$ is even. □

1.3.2 Paths in trees

The first two theorems in this Subsection, Theorems 4 and 5, are the finite and infinite versions, respectively, of the same result: given a set of vertices in a tree, we want to find a collection of pairwise edge-disjoint paths such that each vertex is the end of exactly one of them. This can be done as long as we accept that at most one vertex will not be the end of any path.

Theorem 4. *Let T be a finite tree, S a set of vertices in T. Then T has a collection of pairwise edge-disjoint paths such that each vertex in S, except at most one, is the end of precisely one of the paths.*

Proof. First, suppose that S is a set of $2p$ vertices. We prove the statement by induction on the size of T, i.e. the number $|E(T)|$ of edges of T. If T has size 1 and $p = 2$ or $p = 0$, the statement is true.

Let v be an arbitrary leaf of T, w its neighbor and $e = \{v, w\}$ the edge between v and w.

We will denote by T' the graph obtained from T by removing v and e; note that T' is a tree of size strictly smaller than the size of T.

1. If $v \notin S$, apply induction to T' to find the p required paths.
2. If $v, w \in S$, we choose as one of our paths the edge e. Let $S' = S \setminus \{v, w\}$, which is a set of cardinality $2p - 2$; by applying induction to T', S', we then find $p - 1$ paths in G'. These $p - 1$ paths and the edge e are the p required paths.
3. Otherwise, that is if $v \in S$ but $w \notin S$, we apply induction to $T', S' = S \setminus \{v\} \cup \{w\}$; now $|S'| = 2p$ because $w \notin S$. Among the p paths obtained, denote by a the one that has w as endpoint. We add the edge e to a: this is one of the p required paths. The other $p - 1$ ones are those found in T' (different from a).

If S is a set of odd cardinality, choose an arbitrary element v in S and apply the result to the now even set $S \setminus \{v\}$. □

The infinite version was stated and proven by Thomassen [12]. We give here a sketch of the proof, see the referenced article for the missing details.

Theorem 5. *Let T be a tree and let A be a set of vertices in T. Then T has a collection of pairwise edge-disjoint paths, each joining two vertices in A such that each vertex in A, except possibly one, is the end of precisely one of the paths. If the exceptional vertex in A exists, then it is not the end of any path in the path-collection.*

Proof. If A is finite, the statement is true by Theorem 4.

Now suppose that A is infinite. If A is countable, we enumerate it; if A is uncountable, we consider a well-ordering. Recall that the well-ordering theorem is equivalent to the Axiom of Choice. In order to prove that the statement of the theorem is true for all elements of A we will make use of the induction principle or the transfinite induction principle, according to whether we enumerated A or well-ordered it.

For simplicity, suppose that A can be enumerated, with elements $v_1, v_2, \ldots$

If it exists, we consider a path P_1 joining two distinct vertices v_i, v_j in A such that $T - E(P_1)$ has all vertices of $A_1 = A \setminus \{v_i, v_j\}$ contained in only one component T_1. Repeat the argument: if it exists, we consider a path P_2 contained in T_1 joining two vertices v_p, v_q in A_1 such that all vertices of $A_2 = A_1 \setminus \{v_p, v_q\}$ are contained in only one component T_2 of $T_1 - E(P_2)$. Notice that P_2 is contained in (a component of) $T - E(P_1)$, hence it does not have edges in common with P_1.

In this way we obtain a chain of pairings, ordered by inclusion. By Zorn's lemma, there exists a maximal pairing of vertices. Let A' be the set of vertices in A that do not appear in the maximal pairing; also let T' be the component of T containing the vertices in A'. Notice that T' is a tree and is in general one of possibly many components of the forest obtained by removing from T the edges of the maximal pairing.

By maximality, T', A' satisfy the following property $\mathcal{P}$: it is not possible to find a path P' in T' joining two vertices v_i, v_j in A' such that all vertices of $A' \setminus \{v_i, v_j\}$ are contained in only one connected component of $T' - E(P')$. If A' has one element, this vertex is the exceptional vertex and we have reached our conclusion. If A' is finite, as noted before, the statement is true. So now suppose that A' is infinite. We claim that in this case there is a pairing in T' between all vertices in A', i.e. a collection of pairwise edge-disjoint paths in T' joining vertices in A' such that each vertex in A' is the end of precisely one path in the collection.

Let i be the smallest index such that v_i is in A'; also let P be a path in T' from v_i to another vertex a of A' such that no intermediate vertex of P is in A' and such that as few components of $T' - E(P)$ as possible contain precisely one vertex of A'. Notice that the operation of removing the edges of P from T' (and in general from a graph) leaves a graph with finitely many components (at most $|E(P)| + 1$ components, and exactly $|E(P)| + 1$ components if the edges are removed from a tree), since P has finitely many edges.

We add P to the pairing and delete the edges of P from T'. It is possible to show that in fact no component $T' - E(P)$ contains precisely one vertex of A'. Therefore, by property $\mathcal{P}$, every component of $T' - E(P)$ contains either none or infinitely many vertices of A'.

Consider a component T'' of $T' - E(P)$ that contains infinitely many vertices of A'; let A'' be the set of vertices of A' contained in T''. If T'', A'' have property $\mathcal{P}$, then we repeat the argument applied previously to T', A'. So we can assume that T'', A'' do not have property $\mathcal{P}$. Now let P_1' be a path joining two vertices v_p, v_q in A'' such that all vertices of $A'' \setminus \{v_p, v_q\}$ are contained in only one component of $T'' - E(P_1')$.

Repeat this argument. If we obtain a pairing of A'' we are done; otherwise it is possible to show that, after the deletion of some more paths, we obtain T''', A''' having property $\mathcal{P}$, in which case we repeat the argument previously applied to T', A'.

Repeat the argument. Since at each step we are pairing at least the vertex with the lowest index among those previously unpaired, this will give us the pairing required. □

In the next theorem, given a tree, we find a finite number of collections of paths (in our case 2 collections) such that no vertex of the tree is the endpoint of more than one path in the same collection. Each collection of paths will correspond to a matching by replacing each path with a single edge having the same endpoints as the path, i.e. by ignoring the intermediate vertices of the path. See Section 2.1 for the details.

Theorem 6. *Let T be a tree and v a vertex of T. The edge set of T can be decomposed into pairwise edge-disjoint paths, and these paths can be divided into two classes $\mathcal{P}_1, \mathcal{P}_2$ such that each vertex of T is the end of at most one path in $\mathcal{P}_1$ and at most one path in $\mathcal{P}_2$. Moreover, v is not the end of any path in $\mathcal{P}_1$.*

Proof. First, we pair all edges incident with v (except possibly one), creating paths of length 2 (and possibly one of length 1). These are the blue paths in Figure 1.8 and belong to $\mathcal{P}_2$.

Next, we repeat the argument for each component of $T - v$; for a generic neighbor w of v, we pair all edges incident with w (except the edge vw) creating paths of length 2, and possibly one of length 1. These are the red paths in Figure 1.8 and belong to $\mathcal{P}_1$. Now all edges incident with vertices at distance 1 from v have been covered by paths in either collection.

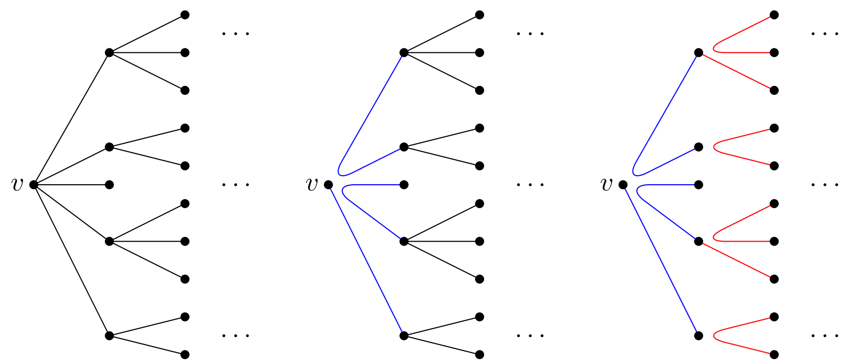

Figure 1.8: Red paths are in $\mathcal{P}_1$, blue paths in $\mathcal{P}_2$.

Repeat the argument for all the remaining vertices of T, alternating between the two colors. $\mathcal{P}_1$ [resp. $\mathcal{P}_2$] is the collection of all red [resp. blue] paths.

Eventually, every edge of T will be colored and therefore contained in one of the two collections, since in a connected graph every vertex has finite distance with any other fixed vertex.

The collections $\mathcal{P}_1, \mathcal{P}_2$ so defined satisfy the conditions of the theorem. In fact, whenever two paths share an endpoint, they have different colors, no vertex is the endpoint of more than two paths, and v is not the endpoint of any path in $\mathcal{P}_1$. □

Finally, we state here a fundamental result linking edge-connectivity and pairwise edge-disjoint spanning trees. See [2] (Corollary 3.5.2).

Theorem 7. *If G is a $2k$-edge-connected multigraph, then G contains at least k pairwise edge-disjoint spanning trees.*

1.3.3 Finite and countable versions of Nash-Williams' decomposition theorem

In the next chapter, we will prove Nash-Williams' cycle decomposition theorem (Theorem 14), which gives a characterization for general graphs that have a decomposition into cycles. We will show that a graph can be decomposed into cycles if and only if it has no finite odd cut.

One implication is easy. Given a decomposition of a graph into cycles, consider a cut. It may intersect none of these cycles (then it is empty), infinitely many of them (then it is also infinite) or a number n of them (then it contains a finite and even number of edges, this number being $\geq 2n$). This shows that any finite cut is necessarily even.

Before studying the general case of Nash-Williams' decomposition theorem, we prove the result for countable graphs.

Theorem 8. *A countable graph has a collection of pairwise edge-disjoint cycles containing all the edges of the graph if and only if the graph has no finite odd cut.*

Proof. We have already shown that for finite graphs, the two conditions given here are equivalent to each other and in fact equivalent to the graph being Eulerian.

Consider now a countable graph G. If it has countably many vertices but only finitely many edges, then each of its connected components is a finite graph, hence we apply the previous argument.

Therefore we can suppose that the graph has countably many edges; we enumerate them: $e_1, e_2, e_3, \ldots$.

Since the graph G is 2-edge-connected, there is a cycle C containing e_1. Remove the edges of C; we obtain a graph G'. We claim that (each component of) G' is again countable and does not contain any finite odd cut. G' is a subgraph of G, hence it is countable. Every cut of G contains an even number of edges contained in C. If the cut is infinite, after the removal of C it becomes an infinite cut in G'; if it is finite (and even), after the removal it becomes a finite and again even cut of G', possibly empty.

Apply the same argument to the edge e_i with the lowest index among the edges that have not been removed. □

Chapter 2

Decompositions into cycles

In this chapter we focus on the problem of decomposing a graph into cycles. The fundamental result is Nash-Williams' cycle-decomposition theorem (Theorem 14 in Section 2.3), which generalizes Theorem 8 to infinite graphs.

We prove the theorem here following the method given by Thomassen in [13]. Curiously, the proof starts with a covering result, rather than a decomposition one. We show that every 2-edge-connected graph can be covered with cycles such that each edge is in at most countably many of them. The proof is detailed in Section 2.1.

Theorem 9 is then used in the proof of Theorem 11, which is a decomposition result into connected, countable and cut-faithful subgraphs; this is done in Section 2.2.

Finally, given a general graph, we can decompose it into countable subgraphs according to Theorem 11 and decompose each of the subgraphs into cycles by the countable version of Nash-Williams' decomposition theorem. This gives a decomposition of a general graph (with no odd cut) into cycles.

2.1 Cycles covering each edge at most countably many times

In Subsection 2.1.1 we define (countable) covering cycle collections, or c.c.c. for short, in Subsection 2.1.2 we show how to pass from a c.c.c. to another of bigger depth and finally in Subsection 2.1.3 we state and prove Theorem 9.

2.1.1 Countable covering cycle collections

In order to cover a graph with cycles, we use the following technique. We define a sequence $(\mathcal{C}_n)_{n\in\mathbb{N}}$ of collections of cycles such that $\mathcal{C}_n$ covers all edges that reach vertices with distance $< n$ from a fixed vertex v (and possibly more edges). Since we will define an increasing sequence, that is we will have $\mathcal{C}_{n-1} \subseteq \mathcal{C}_n$, we can then take the union of all $\mathcal{C}_n$: this is the final collection of cycles covering the graph.

We introduce the terminology *covering cycle collection* to describe this object; properties (i) and (ii) ensure that all edges not "too far" from v are covered as we require.

Let G be a 2-edge-connected graph, v a vertex of G, $n \geq 0$ a natural number. A countable covering cycle collection (of depth n) is a collection $\mathcal{C}_n$ of cycles with the following properties:

(i) Every edge of G is in at most finitely many cycles in $\mathcal{C}_n$. The union of all cycles in $\mathcal{C}_n$ is the edge set of an induced subgraph G_n of G.

(ii) G_n contains all vertices of distance $< n$ to v. It may contain vertices whose distance from v is $\geq n$.

(iii) If M is an external matching joining vertices in the boundary of G_n, then $G_n \cup M$ has a collection $\mathcal{D}$ of cycles such that every edge in M is in at least one cycle in $\mathcal{D}$, and every edge of $G_n \cup M$ is in only finitely many cycles in $\mathcal{D}$.

Since we want to cover a graph with cycles so that each edge of the graph is contained in at most countably many cycles, we add the "in at most finitely many cycles" requirement in property (i). Because of this addition, we need also property (iii), which is more technical and allows us to expand the collections of cycles in a controlled way.

Property (iii) is stated for external matchings, however it holds for general matchings.

(iii*) If M is a matching joining vertices in the boundary of G_n, then $G_n \cup M$ has a collection $\mathcal{D}$ of cycles such that every edge in M is in at least one cycle in $\mathcal{D}$, and every edge of $G_n \cup M$ is in only finitely many cycles in $\mathcal{D}$.

In fact, properties (i) and (iii) together are equivalent to (i) and (iii*). Given a matching M (of a graph G_n), it can be thought as a union $M_E \cup M_I$, where M_E is an external matching and M_I is a matching consisting of edges which were already in G_n. We have that $G_n \cup M = G_n \cup M_E$. Property (iii) applied to M_E gives a collection $\mathcal{D}_E$ of cycles covering each edge of M_E and such that every edge of $G_n \cup M_E$ is in only finitely many cycles in $\mathcal{D}_E$. Now $\mathcal{C}_n \cup \mathcal{D}_E$ is the collection $\mathcal{D}$ required in property (iii*).

A brief remark on the terminology just introduced. According to our definition, a countable c.c.c. is a collection $\mathcal{C}_n$ of cycles in G such that every edge of G is in at most finitely many cycles in $\mathcal{C}_n$. Based on what we have done so far, the name "finite c.c.c." may seem best suited for this object. The name "countable c.c.c." was chosen for two reasons:

- Countable c.c.c.'s are used to prove Theorem 9, according to which every edge of G is covered by *countably* many cycles; in fact, as noted after the theorem, an edge that is covered by some cycles in $\mathcal{C}_n$ may be covered by other cycles belonging to $\mathcal{C}_m$, with $m > n$, and in the final cover of G (which is the union of all $\mathcal{C}_n$'s) there may be countably many cycles passing through a fixed edge. Moreover, in order to prove the theorem, it is not necessary that edges are covered by finitely many cycles in $\mathcal{C}_n$; in fact, it suffices to ask that every edge of G is in at most countably many cycles in $\mathcal{C}_n$. See after Theorem 9 for more details.

- In Section 3.2 we will actually define finite c.c.c.'s, which will also be denoted by $\mathcal{C}_n$; this object is then used to prove Theorem 19, according to which it is possible to cover a 2-edge-connected graph with a collection $\mathcal{C}$ of cycles such that every edge is in at most 7 cycles in $\mathcal{C}$. Countable and finite c.c.c.'s have similar definitions and are both used to find increasing sequences of collections of cycles, whose union gives a cycle cover of the whole graph. The sequences are constructed inductively; we devote Subsection 2.1.2 to show how to pass from a countable c.c.c. $\mathcal{C}_{n-1}$ to a countable c.c.c. $\mathcal{C}_n$ and Subsection 3.2.2 to show the same for finite c.c.c.'s. For this reason the two subsections are called "The first construction" and "The second construction" respectively.

2.1.2 The first construction

Our goal is to define a sequence $(\mathcal{C}_n)_{n\in\mathbb{N}}$ of collections of cycles. The first step is easy: we put $\mathcal{C}_0 = \varnothing$ and $G_0 = \{v\}$.

In the present subsection, we are going to describe how to pass from $\mathcal{C}_{n-1}$ to $\mathcal{C}_n$. Together with the base step, this inductively defines the sequence. See the details in the proof of Theorem 9.

Suppose we have defined a c.c.c. $\mathcal{C}_{n-1}$. We want to construct a c.c.c. $\mathcal{C}_n$ such that $\mathcal{C}_{n-1} \subseteq \mathcal{C}_n$.

We will denote by H a generic connected component of $G - V(G_{n-1})$, as in Figure 2.1.

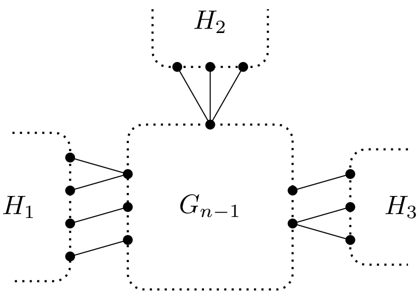

Figure 2.1: H_1, H_2, H_3 are connected components of $G - V(G_{n-1})$.

Step 1: Construction of H-paths and exceptional H-paths

First, we construct two collections of paths, whose elements will be called H-paths and exceptional H-paths. Each path (in either collection) will start and end with an edge between G_{n-1} and H; their intermediate edges will be in H. Moreover the H-paths will be pairwise edge-disjoint. Finally, there will be at most one exceptional H-path for each connected component H of $G - V(G_{n-1})$.

In order to find such paths, we will use the assumption that G is 2-edge-connected.

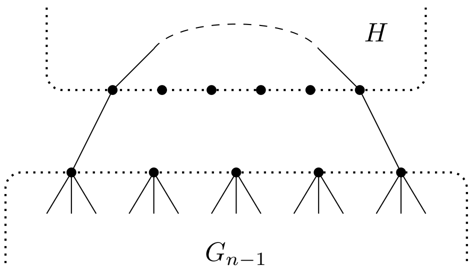

Figure 2.2: We construct paths starting and ending with edges between G_{n-1} and H having all intermediate edges in H.

If a vertex x in H is joined to more than one vertex in G_{n-1}, then we consider a maximal collection of pairwise edge-disjoint paths of length 2 having x as mid-vertex and having their endvertices in G_{n-1}. The collection may contain some double edges, which are considered paths of length 2, between H and G_{n-1}. We remove these paths from G and add them to the collection of H-paths.

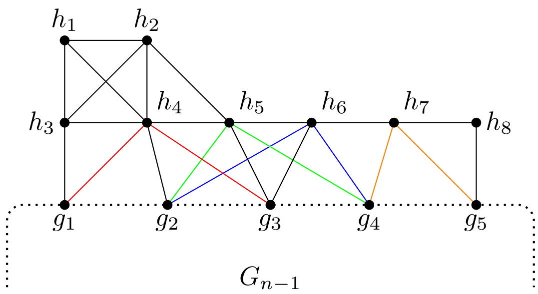

Figure 2.3: We remove the paths $g_1h_4g_3$, $g_2h_5g_4$, $g_2h_6g_4$ and $g_4h_7g_5$.

Now there can be no multiple edge remaining between H and G_{n-1}, otherwise the collection would not be maximal. In particular, for each vertex in H, there is at most one edge connecting it to G_{n-1}. If there are still edges left between G_{n-1} and H, then we let A denote the ends of these edges in H. Let T be a spanning tree of H and apply Theorem 5 to T, A.

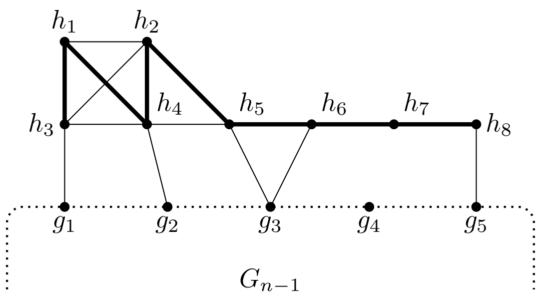

Figure 2.4: A consists of h_3, h_4, h_5, h_6, h_8; the spanning tree T is highlighted.

Theorem 5 gives a collection of pairwise edge-disjoint paths in T, each joining two vertices in A such that each vertex in A, except possibly one, is the end of precisely one of these paths. By appending the corresponding edges between G_{n-1} and H, we can extend these paths to (pairwise edge-disjoint) paths that start and end with an edge between G_{n-1} and H. Some of the extended paths may be cycles; see Figure 2.5. All of the extended paths are removed from G and added to the collection of H-paths.

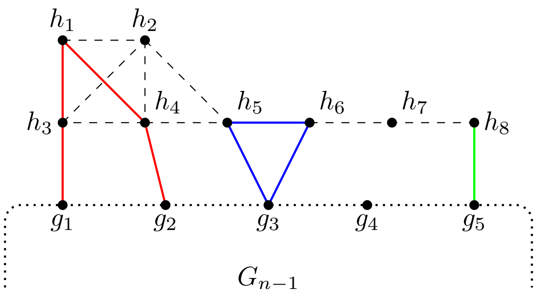

Figure 2.5: $g_1h_3h_1h_4g_2$ and $g_3h_5h_6g_3$ are H-paths (with $g_3h_5h_6g_3$ being a cycle); g_5h_8 is the exceptional H-edge.

Now we describe the case in which Theorem 5 does not pair all vertices in A; in other words, exactly one vertex in A is not the endpoint of any path in the collection of paths given by the theorem. This happens for instance when A is a finite set with odd cardinality.

In this case, exactly one edge between G_{n-1} and H is not used to extend the paths in T and remains not covered by the H-paths described so far. This edge will be called the exceptional H-edge. We cover it with a path (or cycle) starting with the exceptional edge, ending with an edge from H to G_{n-1} and having all intermediate edges in T; this path, which necessarily shares an edge of the form g_ih_j with an H-path (and possibly some more edges in T with other H-paths) is added to the collection of exceptional H-paths. See Figure 2.6.

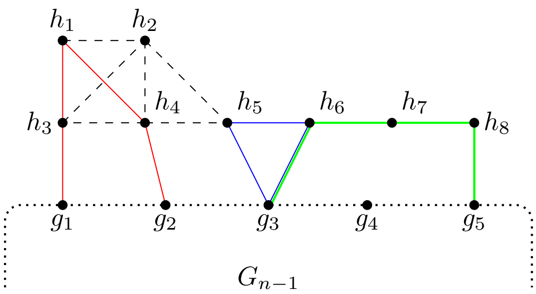

Figure 2.6: The green path $g_5h_8h_7h_6g_3$ is added to the collection of exceptional H-paths.

Now the vertex set of G_n consists of $V(G_{n-1})$, the vertices of the H-paths (for all components H of $G-V(G_{n-1})$), and also the vertices of each exceptional H-path, if it exists. Therefore G_n satisfies (ii), i.e. it contains all vertices of distance $< n$ to v.

Step 2: Arrangement of H-paths and exceptional H-paths into cycles and paths

We now describe how to extend $\mathcal{C}_{n-1}$ to $\mathcal{C}_n$.

First, we consider the collection of all H-paths (not exceptional H-paths) for all components H of $G - V(G_{n-1})$. We take a maximal collection of pairwise edge-disjoint cycles, each of which is the union of H-paths; we add them to $\mathcal{C}_n$ and remove their edges. Then we take a maximal collection of pairwise edge-disjoint cycles, each consisting of some H-paths and some edges in G_{n-1}, add them to $\mathcal{C}_n$ and remove the edges contained in the H-paths.

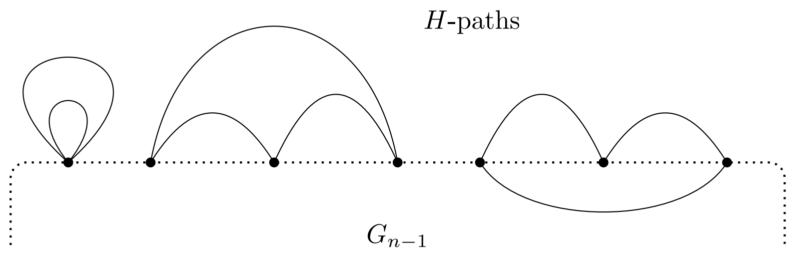

Figure 2.7: We remove maximal collections of cycles obtained as unions of H-paths (and possibly edges in G_{n-1}).

Consider the remaining H-paths. By ignoring their intermediate vertices, we can think of each of them as a single edge joining two vertices in the boundary of G_{n-1}. In other words, we can think of them as a forest F defined on the vertices in the boundary of G_{n-1}.

Apply Theorem 6 to (each component of) F. As a result, we get two collections $\mathcal{P}_1, \mathcal{P}_2$ of paths in F; each of them corresponds in G to a walk, since vertex repetitions may occur. Let $\mathcal{P}_3$ be the collection of all exceptional H-paths.

Consider the collection $\mathcal{P}_1$, whose elements are paths joining vertices in the boundary of G_{n-1}. We think of each of these paths as a single edge, therefore getting a matching M_1. By property (iii*) of G_{n-1} applied to M_1, we get a collection of cycles in $G_{n-1} \cup M_1$.

We replace each edge in M_1 by a union of H-paths and obtain thereby a collection of closed walks. In each such walks there may be repetition of vertices but not edges, and so it can be decomposed into pairwise edge-disjoint cycles. We add these cycles to $\mathcal{C}_n$.

Repeat the argument for $\mathcal{P}_2$ and $\mathcal{P}_3$.

So far we have enlarged $\mathcal{C}_n$ so that it covers all edges in H-paths and exceptional H-paths and no edge is in infinitely many cycles in $\mathcal{C}_n$.

Step 3: Making G_n an induced subgraph

We still need to enlarge $\mathcal{C}_n$ so that the union of its cycles is an induced subgraph. We call residual edge any edge joining vertices of G_n and not contained in any cycle of $\mathcal{C}_n$. Every residual edge joins two vertices of $V(G_n) - V(G_{n-1})$. We need to add enough cycles to $\mathcal{C}_n$ so that every residual edge is covered by at least one of them; the cycles that we will add may have edges in common with cycles already in $\mathcal{C}_n$. Again, we will make sure that after the addition of these new cycles, no edge is in infinitely many cycles of $\mathcal{C}_n$.

First we add to $\mathcal{C}_n$ a maximal collection of pairwise edge-disjoint cycles consisting of residual edges. The remaining residual edges form a forest; by Theorem 6, we decompose it into paths.

Consider one of these paths, say Q, and let x, y be its endpoints; both x, y are on H-paths (or exceptional H-paths), possibly different. By adding edges from these H-paths, we can extend Q to a walk Q' which joins two vertices in the boundary of G_{n-1} and such that the edges of Q' are outside G_{n-1}.

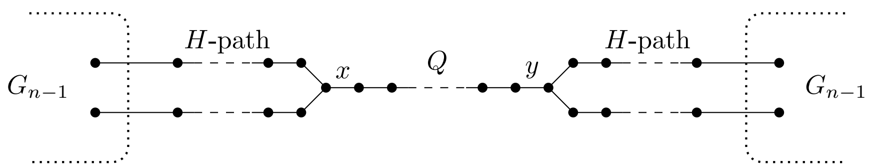

Figure 2.8: Q is extended to a path Q' by taking edges from H-paths.

The walk Q' has no repetition of edges and can therefore be edge-decomposed into cycles and a path with the same ends as Q', which we call Q''. The paths Q'' are not necessarily edge-disjoint; they may share edges in H-paths.

There can be only finitely many paths Q with endpoints on a given H-path (because paths Q have distinct endpoints and H-paths have only finitely many vertices).

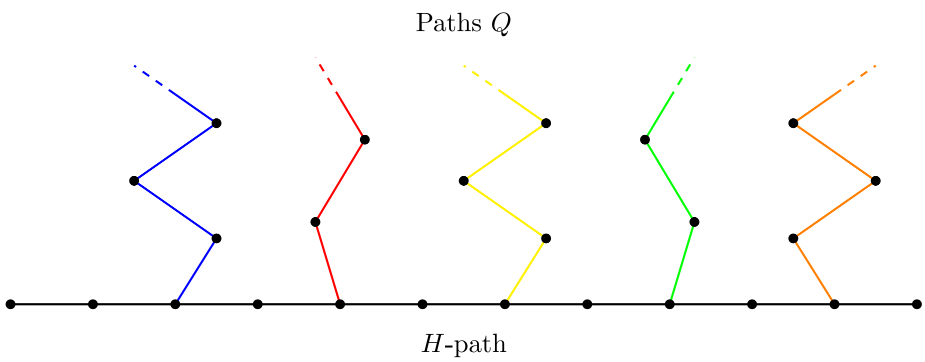

Figure 2.9: There are only finitely many paths Q with endpoints on a given H-path.

Therefore each edge in a fixed H-path is contained in only finitely many paths Q''. We make the paths Q'' pairwise edge-disjoint by replacing each edge by a multiple edge of finite multiplicity.

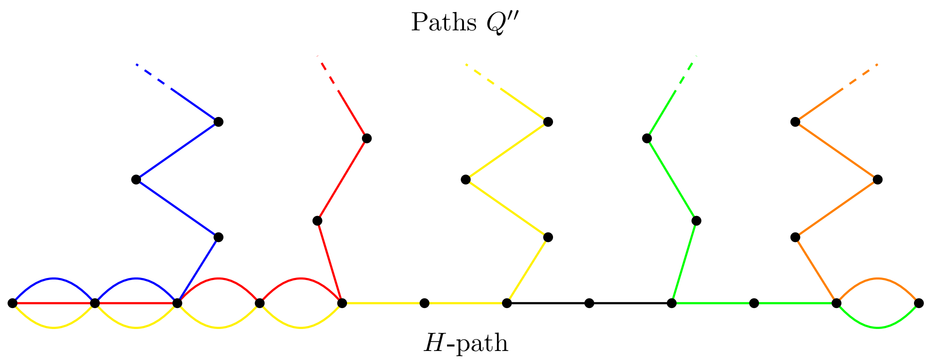

Figure 2.10: Paths Q'' are now pairwise edge-disjoint.

We form a new graph F' whose vertex set is the boundary of G_{n-1}; two vertices u, v are neighbors in F' if there is a path Q'' joining u, v. Remove a maximal collection of pairwise edge-disjoint cycles from F'; if they correspond to cycles in G, we add these cycles to $\mathcal{C}_n$. Since we have created multiple edges, a cycle in F' may consist of a double edge; these do not correspond to cycles in G and will not be added to $\mathcal{C}_n$. The single edges from which these multiple edges originated are in fact already covered by cycles in $\mathcal{C}_n$. Now we may assume that F' is a forest.

We apply Theorem 6 to F' and obtain two collections of paths such that no two paths in the same collection have a common end. Thus we may think of these paths (in either collection) as a matching M consisting of external edges added to the boundary of G_{n-1}.

Let e be an edge in M; it corresponds to a path in F', which in turn corresponds to a walk in G. We can edge-decompose that walk into cycles and a path Q''' in G with the same ends as e. By property (iii) applied to G_{n-1}, M, we find a collection of cycles which contain all paths Q''' and hence all paths Q. We add this collection to $\mathcal{C}_n$ and now G_n is an induced graph, thus satisfying (i).

Now we show that G_n satisfies (iii). Consider a matching M consisting of external edges whose ends are in the boundary of G_n; let e be one of these edges. We repeat for e the same argument used previously for paths Q. Use H-paths to obtain a path with endpoints in G_{n-1}, replace edges of the H-paths with multiple edges of finite multiplicity, consider the graph F' and apply the same argument.

This concludes the construction of $\mathcal{C}_n$.

2.1.3 The cycle covering theorem

Theorem 9. *Let G be a 2-edge-connected graph. Then G has a collection $\mathcal{C}$ of cycles such that every edge of G is in at least one cycle in $\mathcal{C}$ and is in at most countably many cycles in $\mathcal{C}$.*

Proof. Let v be a vertex of G. Our goal is to have a sequence $(\mathcal{C}_n)_{n\in\mathbb{N}}$, where $\mathcal{C}_n$ is a c.c.c. (of depth n) and such that $\mathcal{C}_{n-1} \subseteq \mathcal{C}_n$.

We start by putting $\mathcal{C}_0 = \varnothing$ and $G_0 = \{v\}$. Given $\mathcal{C}_{n-1}$, by the construction shown in the previous section we can define $\mathcal{C}_n$. Therefore by induction we get the required sequence $(\mathcal{C}_n)_{n\in\mathbb{N}}$.

Now define $\mathcal{C} = \mathcal{C}_0 \cup \mathcal{C}_1 \cup \cdots = \bigcup_{i=0}^{\infty} \mathcal{C}_i$, which satisfies the properties of the theorem. □

Recall that in the definition of c.c.c. we asked the following property:

(i) Every edge of G is in at most finitely many cycles in $\mathcal{C}_n$. The union of all cycles in $\mathcal{C}_n$ is the edge set of an induced subgraph G_n of G.

According to Theorem 9, each edge e of G is contained in at most countably many cycles in $\mathcal{C}$. This comes from the fact that e is in at most finitely many cycles in each collection $\mathcal{C}_n$, and we are taking a countable union of these collections. In general, e may be contained in countably many cycles in $\mathcal{C}$ (and not finitely many): for instance, each collection $\mathcal{C}_n$ may bring new cycles not contained in $\mathcal{C}_{n-1}$ and passing through e. This shows that Theorem 9 cannot be refined, with respect to the proof given here.

Because of this, it would be enough to ask that every edge of G is in at most countably many cycles in $\mathcal{C}_n$. In other words, property (i) in the definition of c.c.c. can be replaced with the following condition

(i*) Every edge of G is in at most countably many cycles in $\mathcal{C}_n$. The union of all cycles in $\mathcal{C}_n$ is the edge set of an induced subgraph G_n of G.

The construction shown in Subsection 2.1.2 still works, with minimal adjustments: expressions such as "each edge is in only finitely many cycles in $\mathcal{C}_n$" need to be changed to "each edge is in at most countably many cycles in $\mathcal{C}_n$". However, since property (i*) gives less restrictions than property (i), it may allow for a simplification of the construction in Subsection 2.1.2 and hence of the proof of Theorem 9.

2.2 Cut-faithful countable subgraphs

In Section 2.2, we show that every graph can be decomposed into connected, countable and cut-faithful graphs. First, we prove the result for 2-edge-connected graphs (Theorem 10), then for general graphs (Theorem 11).

Then, we will prove that every k-edge-connected graph has a decomposition into countable and k-edge-connected subgraphs. This is a crucial step that allows us to restrict the proof of Nash-Williams' decomposition theorem (Theorem 14) to the countable case, which we already proved.

These results, together with some extensions of theirs to general cardinals, first appeared in [7]. We will refer in particular to Theorem 11 as Laviolette's decomposition theorem.

Laviolette's proof, given in [7], depends on Nash-Williams' cycle decomposition theorem. Here we follow the approach used by Thomassen in [13]: first we prove Laviolette's results (by using Theorem 9), from which Nash-Williams' decomposition theorem follows as a corollary.

Theorem 10. *Every* 2-edge-connected *graph has an edge-decomposition into connected, countable, cut-faithful graphs.*

In the following proof, for $p, q \in \mathbb{N} \setminus \{0\}$, we will define a decomposition of G into pairwise edge-disjoint 2-edge-connected countable graphs $G_i^{p,q}$, for $i \in I(p,q)$, where $I(p,q)$ is some index set.

Proof. We first describe the case $p = q = 1$, then $p = 1$ and $q \geq 1$, then for all pairs (p,q) with $p \geq 1$ and $q \geq 1$. After having defined all decompositions, we will define the final one into cut-faithful subgraphs.

Decomposition with $p = 1, q = 1$.

Let $\mathcal{C}$ be a collection of cycles of G such that every edge of G is in at least one cycle in $\mathcal{C}$ and at most countably many cycles; $\mathcal{C}$ exists by Theorem 9.

We define a new graph J, whose vertex set is $\mathcal{C}$ (i.e. a vertex in J is a cycle in $\mathcal{C}$); two vertices in J are joined by an edge if the corresponding cycles in $\mathcal{C}$ have at least one edge in common.

In J each vertex has countable degree, hence each component of J is countable. Denote by J_i, for $i \in I(1,1)$, the components of J; then we let $G_i^{1,1}$ be the union of the cycles in $\mathcal{C}$ corresponding to the vertices in J_i, for all $i \in I(1,1)$.

The graphs $G_i^{1,1}$ are pairwise edge-disjoint, 2-edge-connected (because they are unions of cycles), countable subgraphs of G. In general they are not cut-faithful.

Since the graph $G_i^{1,1}$ is countable, it has only countably many finite edge-sets and hence countably many finite minimal cuts; let us denote the finite minimal cuts by $D_1^i, D_2^i, \ldots, D_l^i, \ldots$.

Decomposition with $p = 1, q = 2$.

Consider the cut D_1^i. If it is a cut in G, we do nothing. Now suppose that it is not a cut in G. Since it is a minimal cut in $G_i^{1,1}$, the graph $G_i^{1,1} - D_1^i$ has exactly two components; let us denote by A, B the vertex sets relative to the cut D_1^i and say that D_1^i consists of edges $a_1b_1, a_2b_2, \ldots, a_nb_n$, with $a_1, \ldots, a_n$ in A and $b_1, \ldots, b_n$ in B.

Create two new vertices x, y and connect them as follows. There is an edge between x and each a_j, there is an edge between y and each b_j, for $j \in \{1, \ldots, n\}$; finally there is a single edge xy. Call the resulting graph $\widetilde{G}_i^{1,1}$.

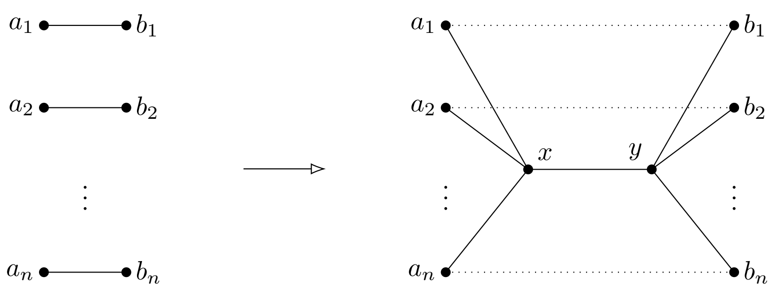

Figure 2.11: The graph $\widetilde{G}_i^{1,1}$ with the two vertices x, y.

For each $i \in I(1,1)$, we have modified the graph $G_i^{1,1}$ into the graph $\widetilde{G}_i^{1,1}$ by operating on only one cut. Since the graphs $G_i^{1,1}$, for $i \in I(1,1)$, are pairwise edge-disjoint, we can perform these operations simultaneously. Moreover, the graphs $\widetilde{G}_i^{1,1}$ so obtained are also pairwise edge-disjoint.

Let $\widetilde{G}$ be the union of all $\widetilde{G}_i^{1,1}$, for $i \in I(1,1)$ (just like G was the union of all $G_i^{1,1}$); since D_1^i is not a cut in G, the graph $\widetilde{G}$ is 2-edge-connected.

We apply Theorem 9 to $\widetilde{G}$ and obtain a collection $\widetilde{\mathcal{C}}$ of cycles such that every edge of $\widetilde{G}$ is in at least one cycle in $\widetilde{\mathcal{C}}$ and at most countably many cycles.

We define a new graph $\widetilde{J}$. The vertex set of $\widetilde{J}$ consists of the union of $\widetilde{\mathcal{C}}$ and the set of all the graphs $\widetilde{G}_i^{1,1}$; in other words, each cycle in $\widetilde{\mathcal{C}}$ and the graph $\widetilde{G}_i^{1,1}$, for each $i \in I(1,1)$, are vertices of $\widetilde{J}$. Two vertices in $\widetilde{J}$ are connected by an edge if the corresponding subgraphs of $\widetilde{G}$ (two cycles of $\widetilde{\mathcal{C}}$ or one cycle of $\widetilde{\mathcal{C}}$ and a graph $\widetilde{G}_i^{1,1}$) have at least one edge in common.

In $\widetilde{J}$ each vertex has countable degree, hence again each of its components is countable. We denote by $\widetilde{J}_j$, for $j \in I(1,2)$, the connected components of $\widetilde{J}$.

For each $j \in I(1,2)$, let $G_j^{1,2}$ be the union of the subgraphs in G (not $\widetilde{G}$) corresponding to the vertices in $\widetilde{J}_j$. This defines a new decomposition of G into pairwise edge-disjoint, 2-edge-connected, countable subgraphs.

The edges of the cut D_1^i are edges of $G_i^{1,1}$ and hence there exists $j \in I(1,2)$ such that D_1^i is contained in $G_j^{1,2}$.

If D_1^i is a cut in G, then it is a cut also in $G_j^{1,2}$. For the remaining indices $i \in I(1,1)$, those such that D_1^i is not a cut in G, we claim that D_1^i is not a cut in $G_j^{1,2}$.

Let $i \in I(1,1)$ be such that D_1^i is not a cut in G; also let C_i be a cycle of $\widetilde{\mathcal{C}}$ that contains the edge of the form xy contained in $\widetilde{G}_i^{1,1}$.

Any graph $G_s^{1,1}$ sharing edges with C_i, for $s \in I(1,1)$, is also in $G_j^{1,2}$; moreover $G_j^{1,2}$ is 2-edge-connected. Therefore C_i can be modified into a cycle in $G_j^{1,2}$ which contains precisely one edge in D_1^i. See Figure 2.12. This implies that D_1^i is not a cut in $G_j^{1,2}$.

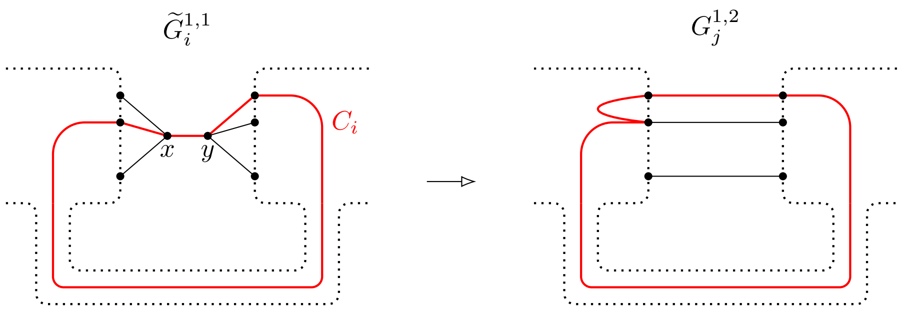

Figure 2.12: C_i is modified into a cycle in $G_j^{1,2}$ containing precisely one edge in D_1^i.

Decompositions with $p = 1, q \geq 3$

So far we have considered the cut D_1^i. Now consider the cut D_2^i, which is also contained in the same $G_j^{1,2}$. If D_2^i is a cut in $G_j^{1,2}$ (and hence a minimal cut) but not a cut in G, we repeat the argument.

In this way we obtain a decomposition of G into pairwise edge-disjoint 2-edge-connected, countable subgraphs $G_k^{1,3}$, for $k \in I(1,3)$. D_2^i is contained in one of these graphs but is not a cut in that graph. Repeat the argument: at step l, we consider D_l^i and obtain a decomposition of G into $G_k^{1,l+1}$, for $k \in I(1, l+1)$.

Decompositions with $p \geq 2, q \geq 1$

Having the decompositions $G_i^{1,q}, i \in I(1,q)$, for all $q \geq 1$, we define a decomposition with $(p,q) = (2,1)$ as follows. Consider an edge e of G; let $i \in I(1,1), j \in I(1,2), k \in I(1,3), \ldots$, be such that e is contained in the graphs $G_i^{1,1}, G_j^{1,2}, G_k^{1,3}, \ldots$, respectively. This is an increasing sequence of graphs; we define their union to be one of the graphs $G_r^{2,1}, r \in I(2,1)$.

p=1	Decomposition $G_i^{1,1}$, for $i \in I(1,1)$, given by Theorem 9
	If D_1^i is not a cut in G, we define a decomposition $G_j^{1,2}$, for $j \in I(1,2)$; now D_1^i is not a cut in $G_j^{1,2}$
	If D_2^i is not a cut in $G_j^{1,2}$, we define a decomposition $G_k^{1,3}$, for $k \in I(1,3)$; now D_2^i is not a cut in $G_k^{1,3}$
	$\vdots$
p=2	Decomposition $G_r^{2,1}$, for $r \in I(2,1)$, defined from the decompositions $G_i^{1,q}, i \in I(1,q)$, for all $q \geq 1$
	If E_1^r is not a cut in G, we define a decomposition $G_s^{2,2}$, for $j \in I(2,2)$; now E_1^r is not a cut in $G_s^{2,2}$
	If E_2^r is not a cut in $G_s^{2,2}$, we define a decomposition $G_t^{2,3}$, for $t \in I(2,3)$; now E_2^r is not a cut in $G_t^{2,3}$
	$\vdots$
p=3	Decomposition $G_u^{3,1}$, for $u \in I(3,1)$, defined from the decompositions $G_i^{2,q}, i \in I(2,q)$, for all $q \geq 1$
	$\vdots$
$\vdots$	$\vdots$

Figure 2.13: We define decompositions $G_i^{p,q}$ for $p = 1$ and all $q \geq 1$, then for $p = 2$ and all $q \geq 1$, and so on.

Each of $D_1^i, D_2^i, \ldots$ is contained in some $G_r^{2,1}, r \in I(2,1)$; moreover, for each

$D_1^i, D_2^i, \ldots$, if it is not a cut in G, then it is also not a cut in $G_r^{2,1}$.

$G_r^{2,1}$ may have new finite minimal cuts; let us denote them by $E_1^r, E_2^r, \ldots$. We repeat the argument and obtain decompositions of G into $G_s^{2,2}, s \in I(2,2)$, then into $G_t^{2,3}, t \in I(2,3)$, and so on.

Having the decompositions $G_i^{2,q}, i \in I(2,q)$, for all $q \geq 1$, we define the decomposition $G_u^{3,1}, u \in I(3,1)$, as before.

Repeat the argument for all $p \geq 1$. Now we have all decompositions of G into $G_i^{p,q}$, for $i \in I(p,q)$, where $p, q \in \mathbb{N} \setminus \{0\}$.

The final decomposition

We define the final decomposition $\mathcal{H}$ of G as follows. Consider an edge e of G; let $i \in I(1,1), r \in I(2,1), u \in I(3,1), \ldots$ be such that e is contained in the graphs $G_i^{1,1}, G_r^{2,1}, G_u^{3,1}, \ldots$, respectively. This is an increasing sequence of graphs; we define their union to be one of the elements of the decomposition $\mathcal{H}$.

Let H be a graph in $\mathcal{H}$; then it is a countable, 2-edge-connected subgraph of G. Since H is the union of an increasing sequence of graphs of the form $G_i^{p,1}$, any finite set of edges in H is contained in a graph $G_i^{p,1}$, for some $p \geq 1$.

We want to show that H is cut-faithful. Let D be a finite minimal cut in H; by definition, this means in particular that $H - D$ is disconnected. We need to show that D is also a cut in G.

Consider all vertices of D contained in one side of $H - D$; since there are finitely many of them, there is a finite connected subgraph of H containing all of them. Call D' the set of edges of this subgraph. Analogously, let D'' be a finite set of edges connecting all vertices of D contained in the other side of $H - D$.

Now $D \cup D' \cup D''$ is a finite set of edges in H, hence it is contained in a subgraph of H of the form $G_i^{p,q}$, for some $p \geq 1$ (and $i \in I(p,1)$). Notice that D is a minimal cut also in $G_i^{p,1}$.

Suppose by contradiction that $G - D$ is connected. Then there is a subgraph $G_j^{p+1,1}$ of H containing D such that $G_j^{p+1,1} - D$ is also connected. Therefore $H - D$ is connected, contradiction. We have that $G - D$ is disconnected, which implies that D is a cut in G.

Therefore $\mathcal{H}$ is a decomposition of G into countable, connected, cut-faithful subgraphs, as required. □

We can easily extend Theorem 10 to general graphs. In fact, any bridge is a countable, connected and cut-faithful subgraph; after their removal, we apply Theorem 10 to each component of the resulting graph.

Theorem 11. *Every graph has an edge-decomposition into connected, countable, cut-faithful graphs.*

Proof. A graph has an edge-decomposition into connected, countable, cut-faithful graphs if and only if each of its connected components does; therefore it suffices to prove the result for connected graphs.

Consider a connected graph. Each of its bridges will be part of the decomposition. In fact the graph •—• with two vertices and one edge joining them is clearly connected and countable; moreover it is cut-faithful if and only if it is a bridge. By removing all the bridges, the remaining graph consists of

2-edge-connected components. We can decompose each of them by Theorem 10. □

We can prove an additional property of decompositions as in Theorem 11.

Theorem 12. *Let G be a graph, $\mathcal{H}$ a decomposition of G into connected, countable, cut-faithful graphs; also let H be an element of $\mathcal{H}$. Then every finite minimal cut of G intersecting H is contained in H.*

Proof. Let D be a finite minimal cut in G intersecting H, i.e. H contains at least one edge of D. Suppose by contradiction that H does not contain all edges of D. Let D' be the set of edges in D which are also in H. Then D' is a cut in H.

Let D'' be a minimal cut in H contained in D'. Then D'' is a proper subset of D and hence not a cut in G (because D is a minimal cut in G). This implies that H is not cut-faithful, contradiction. □

2.3 Nash-Williams' cycle-decomposition theorem

As noted at the beginning of Section 2.2, in order to prove Nash-Williams' decomposition theorem, we will rely on its countable version, Theorem 8. Therefore, we need to prove that it is possible to decompose a 2-edge-connected graph into 2-edge-connected countable subgraphs.

The result, given in a more general form in Theorem 13, depends on Laviolette's decomposition theorem. In fact, the property of the subgraphs given by Theorem 11 of being cut-faithful is what makes preserving the edge-connectivity possible.

Theorem 13. *If k is a natural number, and G is a k-edge-connected graph, then G has an edge-decomposition into countable, k-edge-connected graphs.*

Proof. Decompose G into countable, connected, cut-faithful subgraphs according to Theorem 11. Let H be one the subgraphs; we want to show that H is k-edge-connected.

Suppose by contradiction that H is not k-edge-connected. Then it has a cut, hence also a minimal cut D, which is finite and contains less than k edges. Since H is cut-faithful, D is a cut in G. Therefore G is not k-edge-connected, contradiction. □

Now we can finally prove Nash-Williams' cycle decomposition theorem for general graphs. The result first appeared in [8]; like the rest of this Chapter, we follow here the approach by Thomassen [13] and derive Nash-William's theorem from what we have done so far, in particular from Theorems 8 and 11.

Theorem 14. *Every graph with no odd cut has an edge-decomposition into cycles.*

Proof. Let G be a graph with no odd cut. By Theorem 11, we can decompose G into connected, countable and cut-faithful subgraphs. Let H be one of such graphs.

We claim that H has no finite odd cut.

Suppose by contradiction that H has a finite odd cut C. Then the cut-set of C is the disjoint union of the cut-sets of the finite *minimal* cuts contained in C and since C is odd, at least one of them is odd. Now H is cut-faithful, therefore its finite minimal odd cut is also an odd cut in G, contradiction.

We have shown that H has no finite odd cut. Since it is countable, we can decompose it into cycles by Theorem 8 (Nash-Williams' decomposition theorem for countable graphs).

□

2.4 Summary of the important results

So far we have proven three important theorems:

CCT Cycle covering theorem (Theorem 9): Let G be a 2-edge-connected graph. Then G has a collection $\mathcal{C}$ of cycles such that every edge of G is in at least one cycle in $\mathcal{C}$ and is in at most countably many cycles in $\mathcal{C}$.

LDT Laviolette's decomposition theorem (Theorem 11): Every graph has an edge-decomposition into connected, countable, cut-faithful graphs.

NWDT Nash-Williams' decomposition theorem (Theorem 14): Every graph with no odd cut has an edge-decomposition into cycles.

As we have shown, some of them have easier, but indirect, proofs, done by assuming one of the other theorems:

- Assuming CCT, we proved LDT: this was done in Section 2.2.
- Assuming LDT, we proved NWDT: this was done in Section 2.3.

CCT has an elementary proof assuming LDT.

- Assuming LDT, we show that CCT holds: Given a 2-edge-connected graph, decompose it into connected, countable, cut-faithful (hence 2-edge-connected) subgraphs according to LDT. Then apply the countable version of Nash-Williams' decomposition theorem (Theorem 8) to each subgraph.

Finally, Laviolette [7] proved LDT assuming NWDT.

Chapter 3

Coverings with cycles

A *cycle double cover* of a graph G is a collection of cycles in G such that every edge of G is contained in exactly two cycles in the collection.

A graph has a cycle double cover if and only if each of its connected components has a cycle double cover, so we may suppose that every graph considered here is connected.

If a graph has a cycle double cover, then it is necessarily 2-edge-connected. It has been conjectured that this condition is in fact sufficient. The statement is known as the cycle double cover conjecture, or CDC conjecture for short.

Cycle double cover conjecture. *Every 2-edge-connected graph has a cycle double cover.*

The CDC conjecture is an open problem in graph theory. One of the first published references of the conjecture is a 1973 article by Szekeres [11], although it was formulated for *cubic* graphs; in 1979, Seymour [9] stated the conjecture for all 2-edge-connected graphs. The two formulations are equivalent, see Section 3.1.

In this chapter we present some partial results motivated by the CDC conjecture.

In Section 3.1 we discuss the structure of a potential counterexample to the conjecture.

In Section 3.2 we show that every 2-edge-connected graph has a collection of cycles covering each edge at most 7 times.

In Section 3.3 we give an elementary proof of the fact that every 4-edge-connected graph has a cycle double cover. The same technique gives also the result that Hamiltonian graphs and k-regular, 1-factorable, 2-edge-connected graphs have a cycle double cover.

3.1 Reduction to snarks

Suppose we are given a 2-edge-connected *planar* graph, together with an embedding into the plane. For planar graphs there is a well defined notion of faces.

Since the graph is 2-edge-connected, the boundary of each face, including the external face, is a cycle. Since every edge belongs to exactly two faces, the set of the boundaries of all faces is a cycle double cover of the graph. This shows that the CDC conjecture is true for 2-edge-connected planar graphs. In particular every potential counterexample to the conjecture must be non-planar.

If there exists a counterexample to the CDC conjecture, then there exists also a minimal counterexample. The following theorem (see Jaeger [6]) gives some properties of such minimal counterexample.

Theorem 15. *If G is a minimal counterexample to the CDC conjecture, then:*

(1) G is simple, 3-connected and cubic;
(2) G has no nontrivial 2-cut or 3-cut;
(3) G is not 3-edge-colorable;
(4) G is not planar.

A *snark* is a simple, 2-edge-connected cubic graph with chromatic index 4.

A brief remark on the chromatic index $\chi'(G)$. By Vizing's theorem (Theorem 5.3.2 in [2]), every graph satisfies $\Delta(G) \leq \chi'(G) \leq \Delta(G) + 1$, where $\Delta(G) = \max\{d(v) : v \in V\}$ is the maximum degree of G. Therefore a cubic graph can only have chromatic index 3 or 4.

By Theorem 15 (and the previous observation), a minimal counterexample to the CDC conjecture, if it exists, is a snark. This implies that if the CDC conjecture is true for snarks, then it is true for all (2-edge-connected) graphs.

The smallest snark is the Petersen graph, which admits a cycle double cover.

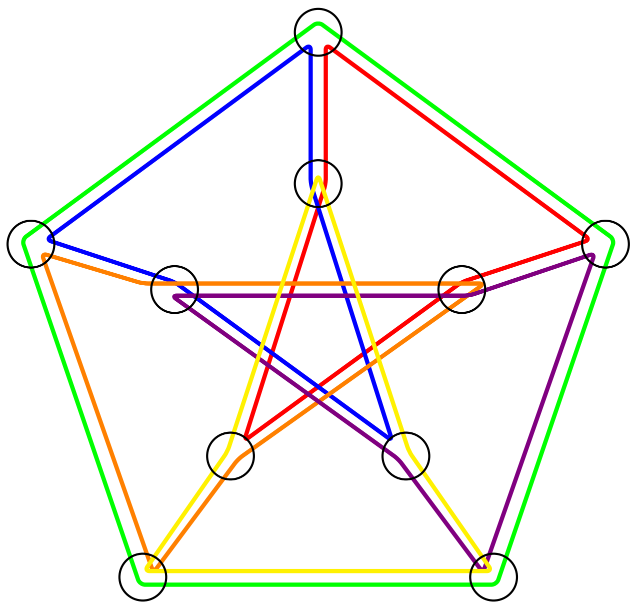

Figure 3.1: A cycle double cover of the Petersen graph.

Huck [4], with the aid of a computer search, showed that the girth, i.e. the length of the smallest cycle, of a potential minimal counterexample to the CDC conjecture is at least 12.

3.2 Cycles covering each edge at most finitely many times

In this Section we are going to show that every 2-edge-connected graph can be covered with cycles such that each edge is in at most 7 of them.

This is a result by Thomassen [13]; we study here the proof given in the referenced article.

The result is very similar to Theorem 9 and so is the proof: in fact, we are going to construct an increasing sequence $(\mathcal{C}_n)_{n\in\mathbb{N}}$ of collections of cycles such that $\mathcal{C}_n$ covers at least all edges that reach vertices with distance $< n$ from a fixed vertex v.

In Section 2.1 we used countable cycle covering collections. Here we need to define a slightly different notion: finite cycle covering collections. Unlike with countable c.c.c.'s, in the case of finite c.c.c.'s an edge that first appears in some of the cycles of $\mathcal{C}_{n-1}$ may be covered by a controlled amount of cycles appearing in $\mathcal{C}_n$, but then does not belong to any cycle in any collection $\mathcal{C}_k$, for $k > n$.

As before, we first give the definition of finite c.c.c.'s (Subsection 3.2.1), then we show how to construct a c.c.c. $\mathcal{C}_n$ given a c.c.c. $\mathcal{C}_{n-1}$ (Subsection 3.2.2) and finally we give the covering result, Theorem 19 in Subsection 3.2.3.

In the proof we will use the following two results.

First, Jaeger [5] proved that every 2-edge-connected graph can covered with three even graphs.

Theorem 16. *Every finite 2-edge-connected graph is the union of three even graphs.*

We say that a vertex v in a graph is *split up into vertices* when we perfom the following operation. v is deleted and replaced by a vertex set V_v; each edge incident with v is replaced with exactly one edge incident with one vertex in V_v.

Thomassen [12] proved that it is possible to split up the vertices of a countable graph while preserving its edge-connectivity.

Theorem 17. *Let k be a natural number, and let G be a countably infinite k-edge-connected graph. Then every vertex of G can be split up into vertices so that the resulting graph is k-edge-connected, and each block of the resulting graph is locally finite.*

3.2.1 Finite covering cycle collections

Let G be a countable, locally finite, 2-edge-connected graph, v a vertex of G, $n \geq 0$ a natural number. A finite covering cycle collection (of depth n) is a *finite* collection $\mathcal{C}_n$ of cycles with the following properties:

(i) The union of all cycles in $\mathcal{C}_n$ is the edge set of an induced subgraph G_n of G. Every edge of G is in at most 7 cycles in $\mathcal{C}_n$.

(ii) G_n contains all vertices of distance $\leq n$ to v.

Given a covering cycle collection $\mathcal{C}_{n-1}$ (of depth $n-1$), we want to construct a c.c.c. $\mathcal{C}_n$ of depth n, with the following additional requirements:

(iii) Every cycle in $\mathcal{C}_{n-1}$ is also in $\mathcal{C}_n$, i.e. $\mathcal{C}_{n-1} \subseteq \mathcal{C}_n$.

(iv) Every edge of G joining two vertices of $V(G_n) \setminus V(G_{n-1})$ is in at most 5 cycles in $\mathcal{C}_n$.

(v) For any two vertices x, y in the boundary of G_n such that $G - V(G_n)$ has a component joined to each of x, y, G_n has a path joining x, y disjoint from G_{n-1}. In other words, for each component H of $G - V(G_n)$, there is a component Q of $G_n - V(G_{n-1})$ such that all edges from H to G_n go to Q.

3.2.2 The second construction

Consider a connected component of $G - V(G_{n-1})$, which we denote by H as before; let S be the set of vertices in G_{n-1} joined to H.

By property (v) applied to H, G_{n-1} (instead of G_n), we have that all vertices of S belong to only one connected component Q of $G_{n-1} - V(G_{n-2})$. By property (iv) applied to G_{n-1}, G_{n-2} (instead of G_n, G_{n-1}), every edge of Q is in at most 5 cycles in $\mathcal{C}_{n-1}$.

We also let U be the set of vertices in H joined to G_{n-1}.

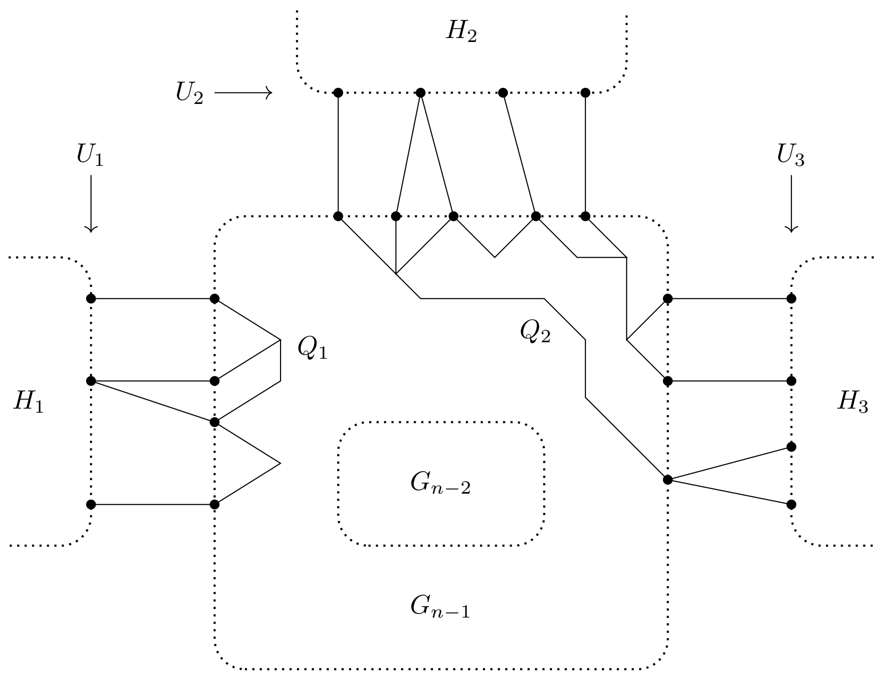

Figure 3.2: Q_1, Q_2 are connected components of $G_{n-1} - V(G_{n-2})$. By property (v), all vertices in S_i belong to one of the components Q_j.

G_{n-1} is finite because it is the subgraph induced by the finite collection $\mathcal{C}_{n-1}$ of cycles. Therefore there are only finitely many edges between G_{n-1} and H. In particular, U is also a finite set.

H has finitely many bridges

An edge e in H is a bridge precisely when e and some edges between H and G_{n-1} form a cut in G.

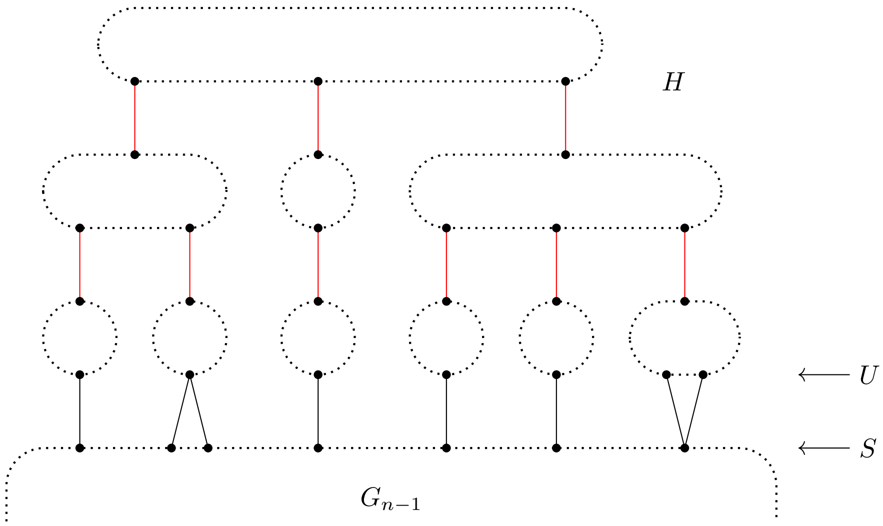

Figure 3.3: Red edges are bridges in H.

There exists a finite connected subgraph of H containing all the vertices in U; for instance, we can take the union of all the paths connecting pairs of such vertices. Every bridge in H is necessarily contained in this subgraph. Therefore there can be only finitely many bridges in H.

Definition of H'

Since H has finitely many bridges, it also has finitely many maximal 2-edge-connected subgraphs.

In each maximal 2-edge-connected subgraph of H, we select a finite 2-edge-connected subgraph which contains:

- all ends of the bridges in H
- and all vertices in U

which are contained in that maximal 2-edge-connected subgraph. Let H' be the union of all these finite subgraphs and of all bridges of H. Then H' is a finite connected subgraph of H; moreover, H, H' have the same bridges.

We may assume that H' is an induced subgraph. In fact, since it is finite, it remains finite even after adding the necessary edges.

By Theorem 16, there exists a collection of cycles in H' covering all its edges (except for the bridges) at least once and at most 3 times. We add all these cycles to $\mathcal{C}_n$.

Construction of the H-paths

We are going to define H-paths and possibly one exceptional H-path, for each connected component H of $G - V(G_{n-1})$, like we did in Section 2.1. All H-paths (including the exceptional ones) will have their endpoints in S and their intermediate vertices in H'. Moreover, they will cover every bridge of H', hence of H, and all the edges between H and S.

First, remove a maximal collection of paths of length 2 consisting of edges between S and H' adjacent to the same vertex in H'. The collection may contain some double edges, which are considered paths of length 2, between S and H'. Moreover, there can be no multiple edge remaining between S and H', otherwise the collection would not be maximal. Let U' be the set of endpoints in H' of all the remaining edges between S and H'.

Let T be a spanning tree of H' and apply Theorem 4 to T, U'. We obtain a collection of pairwise edge-disjoint paths in H' such that each vertex in U' (except possibly one) is the end of precisely one of such paths. Extend each path to a path (or possibly a cycle) having both endpoints in S by adding the corresponding edges between H' and S. These extended paths are all pairwise edge-disjoint H-paths.

If the exceptional vertex exists, there is one exceptional edge between H' and S that has not been covered by the H-paths. Then we find a path in T from the exceptional vertex to another vertex in U'. We extend this path by adding the exceptional edge and the corresponding edge between H' and S adjacent to the other endpoint. The extended path is the exceptional H-path and, if it exists, it necessarily shares edges with one of the H-paths.

Covering the H-paths

Recall that we denoted by Q the connected component of $G_{n-1} - V(G_{n-2})$ that contains all vertices in S. Consider all H-paths and exceptional H-paths, for all components H of $G - V(G_{n-1})$ which are joined to Q.

By ignoring their intermediate vertices, each H-path (not the exceptional H-paths) can be thought as an external edge added to Q. Let Q' be the graph obtained from Q by adding the external edges; also let T_Q be a spanning tree of Q.

We claim that by removing edges from T_Q, it is possible to transform Q' into an even graph Q''. Apply Theorem 4 to T_Q, where S is the set of vertices of odd degree in Q'; then remove the edges of the paths given by the Theorem. The remaining edges of Q' form an even graph Q''. Decompose Q'' into pairwise edge-disjoint cycles; the corresponding cycles in G cover all H-paths and are added to $\mathcal{C}_n$.

All edges of $G_n - V(G_{n-1})$ are covered at most 4 times: they belong to at most 3 cycles obtained from Theorem 16 and at most one cycle coming from the decomposition of Q''.

All edges of $G_{n-1} - V(G_{n-2})$ are covered at most 6 times: they belong to at most 5 cycles in $\mathcal{C}_{n-1}$ by assumption and at most one cycle coming from the decomposition of Q''.

Now we cover the exceptional H-paths. Each of them is again thought as an external edge added to Q. By repeating the argument used for H-paths, we can find a collection of pairwise edge-disjoint covering the exceptional H-paths. These cycles are also added to $\mathcal{C}_n$.

Therefore all edges of $G_n - V(G_{n-1})$ are covered at most 5 times and all edges of $G_{n-1} - V(G_{n-2})$ are covered at most 7 times.

$\mathcal{C}_n$ satisfies properties (ii), (iii) and (iv). Since H' is induced, G_n is an induced subgraph as well, hence also property (i) is satisfied.

We want to show that property (v) holds. Consider a connected component J of $G - V(G_n)$. All its edges go to only one component H of $G - V(G_{n-1})$, otherwise there would be two (or more) distinct connected components of $G - V(G_{n-1})$ which are connected to each other, contradiction.

Moreover, we claim that all edges between J and G_n go to the subgraph H' relative to that component H, and not to G_{n-1}. In fact, all edges between G_{n-1} and H (i.e. between the sets S, U) are covered by cycles in $\mathcal{C}_n$, hence they belong to G_n. Notice that we only covered edges in H', therefore every edge in H but outside H' is not in G_n.

We showed that every edge between J and G_n goes to H', which is a connected subgraph, hence property (v) holds for G_n.

3.2.3 The covering result

First we prove the result for countable and locally finite graphs, then for general (2-edge-connected) graphs.

Theorem 18. *Let G be a countable, locally finite, 2-edge-connected graph. Then G has a collection $\mathcal{C}$ of cycles such that every edge of G is in at least one cycle in $\mathcal{C}$ and in at most 7 cycles in $\mathcal{C}$.*

Proof. Let v be a vertex of G. Our goal is to have a sequence $(\mathcal{C}_n)_{n\in\mathbb{N}}$, where $\mathcal{C}_n$ is a finite c.c.c. (of depth n), satisfying all the properties (i) to (v).

We start by putting $\mathcal{C}_0 = \varnothing$ and $G_0 = \{v\}$. Given $\mathcal{C}_{n-1}$, by the construction shown in the previous section we can define $\mathcal{C}_n$. Therefore by induction we get the required sequence $(\mathcal{C}_n)_{n\in\mathbb{N}}$.

Now define $\mathcal{C} = \mathcal{C}_0 \cup \mathcal{C}_1 \cup \cdots = \bigcup_{i=0}^{\infty} \mathcal{C}_i$, which satisfies the properties of the theorem. □

The covering result:

Theorem 19. *Let G be a 2-edge-connected graph. Then G has a collection $\mathcal{C}$ of cycles such that every edge of G is in at least one cycle in $\mathcal{C}$ and in at most 7 cycles in $\mathcal{C}$.*

Proof. By Theorem 13 we can decompose G into countable, 2-edge-connected subgraphs. We will cover each of these subgraphs with cycles. Hence we can suppose that G is countable (and 2-edge-connected).

We split every vertex of G up into vertices according to Theorem 17; now each block of the resulting graph is locally finite and the edge-connectivity is preserved. Apply Theorem 18. □

3.2.4 Directed graphs

There is a result analogous to Theorem 19 for directed graphs.

Theorem 20. *Let G be a 2-edge-connected graph. Then the edges of G can be oriented so that the resulting directed graph has a collection of directed cycles such that each edge is in at least one and finitely many directed cycles in the collection.*

Proof. It follows from Theorem 11 in [12]. □

3.3 Graphs with a cycle double cover

In Section 3.3 all graphs are finite.

The proof we are going to present here is elementary and is based on the following ideas:

- If a graph can be written as the union of two Eulerian subgraphs, then it has a cycle double cover; the proof of this result, Theorem 21, passes through the symmetric difference of the two subgraphs.
- Given a spanning tree T of a graph G, it is possible to find an Eulerian subgraph of G that covers at least every edge not in T. This is done by finding pairwise edge-disjoint paths in T pairing vertices of odd degree in G.
- Since we need two Eulerian subgraphs, whose union is the whole graph, we need to find two edge-disjoint spanning trees; this is possible if the graph is 4-edge-connected. See Theorem 22.
- A Hamiltonian cycle, which is an Eulerian subgraph, may play the same role as the spanning tree T. See Theorem 23.

Singularly, each of these ideas is very easy and has been used extensively, sometimes without the need of putting a reference. It is hard to establish whether the whole proof is original; in fact, we could not find a published reference explicitly mentioning this method and applying it to 4-edge-connected graphs and to Hamiltonian graphs.

First, we state and prove Theorem 21, which gives a sufficient condition for having a cycle double cover.

Theorem 21. *If a graph is the union of two Eulerian subgraphs, then it has a cycle double cover.*

Proof. Let G_1, G_2 be two Eulerian subgraphs whose union is a graph G. By Theorem 3, the symmetric difference $G_1 \triangle G_2$ is also an Eulerian subgraph of G. Together, the subgraphs $G_1, G_2, G_1 \triangle G_2$ cover each edge of G exactly twice. By decomposing each of these three subgraphs into pairwise edge-disjoint cycles, we get a cycle double cover of G. □

3.3.1 Every 4-edge-connected graph has a cycle double cover

Theorem 22. *Every 4-edge-connected graph has a cycle double cover.*

Proof. Let G be a 4-edge-connected graph. We want to show that G can be written as the union of two Eulerian subgraphs; then the result follows by Theorem 21.

Let S be the set of vertices of G with odd degree. By Theorem 7 (with $k = 2$), G has 2 edge-disjoint spanning trees, say T_1, T_2.

By Theorem 4 applied to T_1, S, we can find pairwise edge-disjoint paths $p_1, p_2, \dots, p_n$ in T_1 such that each vertex in S is the end of precisely one of the paths. Let G_1 be the subgraph of G containing all edges of G except the edges of the paths $p_1, p_2, \dots, p_n$.

Similarly, we can find pairwise edge-disjoint paths $p'_1, p'_2, \dots, p'_n$ in T_2 such that each vertex in S is the end of precisely one of the paths. We let G_2

be the subgraph of G containing all edges of G except the edges of the paths $p'_1, p'_2, \ldots, p'_n$.

The union of G_1, G_2 is G (because the trees T_1, T_2 are pairwise edge-disjoint and hence also the paths $p_1, p_2, \ldots, p_n$ and $p'_1, p'_2, \ldots, p'_n$ are all pairwise edge-disjoint); moreover, G_1, G_2 are Eulerian subgraphs, since the paths removed to obtain each subgraph join precisely the vertices of G of odd degree. Apply Theorem 21. □

3.3.2 Every Hamiltonian graph has a cycle double cover

Theorem 23. *Every Hamiltonian graph has a cycle double cover.*

Proof. Let G be a Hamiltonian graph. As in the proof of Theorem 22, we want to show that G can be written as the union of two Eulerian subgraphs.

Let w be a Hamiltonian tour of G; w is one of the two Eulerian subgraphs, say G_1.

Starting from an arbitrary vertex v of G, let $v_1, v_2, \ldots, v_{2k}$ be the vertices of odd degree in G as they appear in w. If $k = 0$, we take G as the subgraph G_2. Otherwise, we consider the paths p_1 from v_1 to v_2 (in w), p_2 from v_3 to v_4 (in w), ..., p_k from v_{2k-1} to v_{2k}, all contained in w. Notice that these paths are all pairwise edge-disjoint.

G_2 is defined to be the subgraph containing all edges of G except the edges of the paths $p_1, p_2, \ldots, p_k$. G_2 is Eulerian because each of the paths $p_1, p_2, \ldots, p_k$ joins vertices of odd degree in G, with each vertex of odd degree being the endpoint of exactly one of the paths.

As before, G_1, G_2 are Eulerian subgraphs whose union is G, hence the result follows by Theorem 21. □

3.3.3 1-factorable graphs

Let G be a graph. A *k-factor* of G is a spanning k-regular subgraph G. A *k-factorization* of G is a decomposition of G into pairwise edge-disjoint k-factors. Finally, G is said to be *k-factorable* if it has a k-factorization.

A 1-factor is a perfect matching, i.e. a matching that contains every vertex of the graph. If a graph is 1-factorable (and the 1-factorization contains k 1-factors), then each of the 1-factors increases by 1 the degree of every vertex of the graph; therefore the graph is k-regular; the converse is not true in general.

A 1-factorization of a graph (into k 1-factors) gives an edge-coloring of the graph with k colors. Conversely, if a k-regular graph can be k-edge-colored, hence it has chromatic index k, then it is 1-factorable (and again into k 1-factors). In fact, every vertex has degree k, hence it must be incident with an edge of each color.

Theorem 24. *Let G be a k-regular, 1-factorable and 2-edge-connected graph, with $k \geq 2$. Then G has a cycle double cover.*

Proof. If k is even, then G is Eulerian; decompose it into pairwise edge-disjoint cycles and take each of them twice.

Suppose that k is odd and let $S_1, S_2, \ldots, S_k$ be the k 1-factors of G. Then $G_1 = G - S_1, G_2 = G - S_2$ are two Eulerian subgraphs whose union is G. The result follows from Theorem 21. □

It is easy to characterize k-regular graph with $k \in \{0, 1, 2\}$: 0-regular graphs consist of only isolated vertices, 1-regular graphs consist of isolated edges, 2-regular graphs consist of isolated cycles.

The following graphs have been proven to be 1-factorable:

- any k-regular graph with chromatic index k (see before Theorem 24);
- any k-regular bipartite graph (Corollary 2.1.3 in [2]);
- any complete graph on an even number of vertices (Theorem 9.1 in [3]);
- any k-regular graph on $2n$ vertices, with $k \in \{2n-1, 2n-2, 2n-3, 2n-4, 2n-5\}$, or with $k \geq \frac{12n}{7}$ (Chetwynd & Hilton [1]).

3.3.4 Final remarks

The converses of Theorems 22 and 23 are not true, with the Petersen graph being a counterexample for both. The Petersen graph does not even satisfy Theorem 21.

However, we can give a generalization of Theorem 21. Instead of having just 2 Eulerian subgraphs, we consider a number n of Eulerian subgraphs such that each edge of the graph is contained in one or two of them. The proof is analogous to the proof of Theorem 21.

Theorem 25. *Suppose that a graph G has a collection of Eulerian subgraphs such that each edge of G is contained in one or two of these subgraphs. Then G has a cycle double cover.*

Proof. Let $G_1, G_2, \ldots, G_n$ be such Eulerian subgraphs. Let G_0 be the subgraph of G consisting of all edges of G contained in exactly one of the subgraphs $G_1, G_2, \ldots, G_n$; then G_0 is also Eulerian. Decompose each of the subgraphs $G_0, G_1, G_2, \ldots, G_n$ into pairwise edge-disjoint cycles. □

The converse of Theorem 25 is also true; in other words, the following two facts are equivalent for a graph G:

- G has a collection of even subgraphs such that each edge of G is contained in one or two of these subgraphs;
- G has a cycle double cover.

Therefore the cycle double cover conjecture can be equivalently stated as follows:

Cycle double cover conjecture. *Every 2-edge-connected graph G has a collection of even subgraphs such that each edge of G is contained in one or two of these subgraphs.*

Bibliography

[1] Amanda G Chetwynd and Anthony J W Hilton. Regular graphs of high degree are 1-factorizable. *Proceedings of the London Mathematical Society*, 3(2):193–206, 1985.

[2] Reinhard Diestel. *Graph Theory*. Springer-Verlag, 1997.

[3] Frank Harary. *Graph theory*. Addison-Wesley, Reading, MA, 1969.

[4] Andreas Huck. Reducible configurations for the cycle double cover conjecture. *Discrete applied mathematics*, 99(1-3):71–90, 2000.

[5] François Jaeger. Flows and generalized coloring theorems in graphs. *Journal of Combinatorial Theory, Series B*, 26(2):205–216, 1979.

[6] François Jaeger. A survey of the cycle double cover conjecture. In *North-Holland Mathematics Studies*, volume 115, pages 1–12. Elsevier, 1985.

[7] François Laviolette. Decompositions of infinite graphs: Part i - bond-faithful decompositions. *Journal of Combinatorial Theory, Series B*, 94(2):259–277, 2005.

[8] Crispin St J A Nash-Williams. Decomposition of graphs into closed and endless chains. *Proceedings of the London Mathematical Society*, 3(1):221–238, 1960.

[9] Paul D Seymour. Sums of circuits. *Graph theory and related topics*, 1:341–355, 1979.

[10] Lajos Soukup. Infinite combinatorics: from finite to infinite. In *Horizons of combinatorics*, pages 189–213. Springer, 2008.

[11] George Szekeres. Polyhedral decompositions of cubic graphs. *Bulletin of the Australian Mathematical Society*, 8(3):367–387, 1973.

[12] Carsten Thomassen. Orientations of infinite graphs with prescribed edge-connectivity. *Combinatorica*, 36(5):601–621, 2016.

[13] Carsten Thomassen. Nash-Williams cycle-decomposition theorem. *Combinatorica*, 37(5):1027–1037, 2017.