

Canonical graph decompositions via coverings

Diestel, Reinhard; Jacobs, Raphael W.; Knappe, Paul; Kurkofka, Jan

DOI:

[10.48550/arXiv.2207.04855](https://doi.org/10.48550/arXiv.2207.04855)

License:

Other (please provide link to licence statement)

Document Version

Other version

Citation for published version (Harvard):

Diestel, R, Jacobs, RW, Knappe, P & Kurkofka, J 2023 'Canonical graph decompositions via coverings' arXiv.
<https://doi.org/10.48550/arXiv.2207.04855>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

CANONICAL GRAPH DECOMPOSITIONS VIA COVERINGS

REINHARD DIESTEL, RAPHAEL W. JACOBS, PAUL KNAPPE, AND JAN KURKOFKA

ABSTRACT. We present a canonical way to decompose finite graphs into highly connected local parts. The decomposition depends only on an integer parameter whose choice sets the intended degree of locality. The global structure of the graph, as determined by the relative position of these parts, is described by a coarser *model*. This is a simpler graph determined by the decomposition, not imposed.

The model and decomposition are obtained as projections of the tangle-tree structure of a covering of the given graph that reflects its local structure while unfolding its global structure. In this way, the tangle theory from graph minors is brought to bear canonically on arbitrary graphs, which need not be tree-like.

Our theorem extends to locally finite quasi-transitive graphs, and in particular to locally finite Cayley graphs. It thereby yields a canonical decomposition of finitely generated groups into local parts, whose relative structure is displayed by a graph.

1. INTRODUCTION

The question to what extent graph invariants—the chromatic number, say, or connectivity—are of local or global character, and how their local and global aspects interact, drives much of the research in graph theory both structural and extremal [36]. In this paper we offer such studies a possible formal basis.

We show that every finite graph G decomposes canonically into local parts which, between them, form a global structure displayed by a simpler graph H . Both H and the decomposition $(H, (G_h)_{h \in H})$ of G into parts indexed by the nodes of H are unique once we have set an integer parameter $r > 0$ to define our desired threshold between ‘local’ and ‘global’: cycles in G of length $\leq r$ are considered as local, and are reflected inside the parts G_h of the decomposition, while the cycles of H reflect only global cycles of G , longer cycles that are not generated by the short ones.




FIGURE 1. The global structure of G is displayed by a cycle H . Its local parts are K^5 s.

Our main result reads as follows:

Theorem 1. *Let G be any finite graph, and $r > 0$ an integer. Then G has a unique canonical decomposition modelled on another finite graph $H = H(G, r)$ that displays its r -global structure.*

See Sections 1.1–1.3 below for a summary of the definitions needed here. More details are given in [Definition 5.4](#) and the text following it.

2020 *Mathematics Subject Classification.* 05C83, 57M15, 05C25, 05C40, 05C63, 05C38.

Key words and phrases. Graph, local, global, connectivity, covering space, tangle, canonical tree-decomposition, Cayley graph, finitely generated group, accessible group.

Since $H = H(G, r)$ and the associated decomposition of G are canonical, they are invariants of G , and we can study how they interact with other graph invariants as r ranges between 1 and $|G|$. For example, we may study how the chromatic number, connectivity, or edge density of G is reflected in H , and is in this sense global, or is reflected in the parts G_h of the decomposition of G , and is in that sense local.

Theorem 1 extends to connected Cayley graphs even when these are only locally finite. It thus implies a canonical decomposition theorem for finitely generated groups into local parts, whose relative structure is displayed by a graph:

Theorem 2. *Let Γ be a group given with a finite set S of generators. For every integer $r > 0$, the Cayley graph $\text{Cay}(\Gamma, S)$ has a unique canonical decomposition modelled on a graph $H = H(\Gamma, S, r)$ that displays the r -global structure of the group Γ as presented by S .*

Our results have further group-theoretical implications. We shall see that finitely generated groups Γ extend to finitely presented groups Γ_r that describe their r -local structure, the aspect of their structure that complements their r -global structure given by **Theorem 2**. Similarly, we shall see that the r -local structure of any finite graph G that complements its r -global structure as described in **Theorem 1** is displayed by a graph G_r modelled on the Cayley graph of a finitely presented group. This graph G_r is typically infinite. It describes the local structure of G by way of a canonical covering $G_r \rightarrow G$. See **Section 4** for details. More group-theoretical applications of this approach have been obtained by Carmesin [3].

The main idea for the construction of $H = H(G, r)$ was inspired by the aims of [4]. The idea is to apply the theory of tangles from graph minor theory, though not to G itself, but to the covering space G_r of G that reflects its short cycles, those of length $\leq r$, while unfolding in a tree-like way any longer cycles not generated by the short ones.¹ This tree-likeness of this graph G_r enables us to use tangles to structure G_r as customary in graph minor theory: by canonical tangle-distinguishing tree-decompositions. Projecting this structure back to G defines H and the decomposition of G modelled on H .

Our original proof of **Theorem 1** also used group-theoretic tools in an essential way. In order to apply tangle-distinguishing tree-decompositions to our covering graph G_r , we needed that G_r resembles a Cayley graph of some finitely presented group that we could decompose using Stallings' theorem [37], and then use [40] to analyse its ends. We formalised this resemblance by showing that G_r had a decomposition into bounded-sized finite parts modelled on a Cayley graph of its group of deck transformations over G . This group is finitely presented not only when G is finite (**Theorem 1**), but also when it is the Cayley graph of an arbitrary finitely generated group (**Theorem 2**). We have since been able to replace these group-theoretical tools by purely combinatorial arguments, which is how we present our proof here.

In the remainder of this introduction we indicate just enough background to be able to make the statement of **Theorem 1** precise, and to give a brief sketch of its proof.

1.1. Tree-decompositions vs. graph-decompositions. Early in their seminal work on graph minors, Robertson and Seymour [34] introduced *tree-decompositions* of graphs as a way to display their global structure when this is tree-like. However, not all graphs are globally tree-like. For those that are not, tree-decompositions are a poor fit to display their global structure.

Yet, conversely, few graphs are totally homogeneous at all levels of local or global focus: most have some clusters that are particularly highly connected. These clusters then form a global structure between

¹ The use of coverings to analyse graph connectivity from a local/global perspective goes back to Amit and Linial [1].

them, even if this is not tree-like. **Theorem 1** makes this global structure visible in the form of a graph H that need not be a tree. Indeed, H will change with the graph G being decomposed and with the locality parameter r we choose: the global structure of G is *found* by our theorem, not imposed as in the case of tree-decompositions.

Let us make this precise. Let H be a graph. A *decomposition* of a graph G *modelled on* H , or an *H -decomposition* of G , is a pair $(H, (G_h)_{h \in H})$ consisting of H and a family $(G_h)_{h \in H}$ of subgraphs of G , the *parts* of this decomposition, that are associated with the nodes h of H in such a way that

- $\bigcup_{h \in H} G_h = G$;
- for every vertex $v \in G$, the subgraph of H induced by $\{h \in H \mid v \in G_h\}$ is connected.

The decompositions of a graph that are modelled on trees are thus precisely its tree-decompositions [10]. The two conditions ensure that H is not merely the index set of a cover of the vertex set of G but reflects, as a graph, the coarse overall structure of G .

Our decompositions and their models $H = H(G, r)$ are *canonical* in that they commute with graph isomorphisms: if $(H, (G_h)_{h \in H})$ and $(H', (G'_{h'})_{h' \in H'})$ are decompositions that witness **Theorem 1** for graphs G and G' , then any graph isomorphism $\sigma: G \rightarrow G'$ maps the parts G_h of G to the parts $G'_{h'}$ of G' in such a way that $h \mapsto h'$ is a graph isomorphism $H \rightarrow H'$. In particular, the graph H on which **Theorem 1** models a given graph G is unique, up to isomorphism, for every choice of r .

1.2. Construction of $H = H(G, r)$. The construction of $H = H(G, r)$ is not hard to describe. It works as follows.

Given G and r , let $\pi_1^r(G, v_0)$ denote the subgroup of the fundamental group $\pi_1(G, v_0)$ of G , based at a vertex v_0 , that is generated by the elements represented by a walk in G from v_0 to a cycle of length at most r , round it, and back along the access path. This $\pi_1^r(G, v_0)$ is a normal subgroup of the fundamental group of G , independent of the choice of basepoint; let $p_r: G_r \rightarrow G$ denote the normal covering of G with $\pi_1^r(G, v_0)$ as characteristic subgroup.

By our choice of r , the cover G_r reflects all the short cycles of G , as well as those of its longer cycles that are generated in $\pi_1(G)$ by the short ones. The other longer cycles of G are usually unfolded to 2-way infinite paths, or *double rays*. (One can construct examples where G_r covers G with finitely many sheets, but those are rare.) Thus, G_r mirrors G in its ‘ r -local’ aspects, but not in its ‘ r -global’ aspects, where it is simply tree-like. This is why we call G_r the *r -local* cover of G .

As a consequence of this tree-likeness of the global aspects of G_r , tree-decompositions will be a better fit for G_r than they were for G , whose global aspects could include long cycles that would not fittingly be captured by tree-decompositions. We exploit this by applying to G_r , not to G , the tree-decompositions that represent the state of the art from the theory of graph minors: those that distinguish all the maximal blocks, tangles and ends of G efficiently, and are canonical in the sense described earlier [13, 19]. We fix one such canonical tree-decomposition that is particularly natural, and use it to define the r -global structure of G . More on this in **Section 1.3**.

Since our tree-decomposition $(T_r, (V_t)_{t \in T_r})$ of G_r is canonical, the group \mathcal{D}_r of deck-transformations of G_r over G acts on its model, the tree T_r . The canonical model H for our r -local decomposition of G is then obtained as the orbit space T_r/\mathcal{D}_r , which is a graph. The parts G_h of this decomposition are the projections to G of the parts V_t of our canonical tree-decomposition of G_r .

1.3. **Details of Theorem 1.** In order to make [Theorem 1](#) precise we have to define what it means for a graph H to ‘display the r -global structure’ of a given graph G . We shall first give a more formal definition of the ‘ r -global structure’ of G , and then define what it means that H displays it. Let $r > 1$ be given.

Consider the r -local cover G_r of G from [Section 1.2](#). This is unique up to isomorphisms (of coverings of G), and it is clearly a graph. The r -global structure of G is formally defined as a particular canonical *tree of tangles* of G_r along with the group \mathcal{D}_r of deck transformations of G_r over G . Such a ‘tree of tangles’ is not a graph. It is a central notion in the theory of graph minors; let us indicate briefly what it means.

Our covering graph G_r may have highly cohesive regions, or ‘clusters’. These will include lifts of the local clusters in G that our H -decomposition of G is meant to distinguish, in that they should live in distinct parts G_h . There are various notions of such clusters in graph theory. The most general of these are known as *blocks* [7,8] and as *tangles* [35]. When G_r is infinite, as it usually will be, it will also have *ends*. Ends can be formalised as infinite tangles, and all blocks except some irrelevant small ones also induce tangles. We shall therefore use the term ‘tangle’ from now on to cover all three of these: blocks, ends, and traditional tangles in the sense of Robertson and Seymour [34]. Details will be given in [Section 5.1](#).

All the tangles of our r -local cover G_r of G can be ‘distinguished’ by a nested set N of separations of G_r , in the following sense. Given any separation of G_r , of order k say, every tangle ‘of order $> k$ ’ in G_r will lie essentially on one of the two sides of this separation, and thereby *orient* it towards that side. A set N of separations is said to *distinguish* the tangles of G_r if for any two of them that orient some separation of G_r differently there is such a separation even in N . If N is nested and even contains, for every pair of tangles, such a separation of minimum order (among all the separations of G_r that distinguish this pair of tangles), we say that N distinguishes the tangles of G_r *efficiently* and call it a *tree of tangles* for G_r .

While G_r can have many trees of tangles, one of these stands out. Let D be the set of all separations of G_r that distinguish some pair of tangles efficiently. Now given any pair of tangles that can be distinguished by some separation of G_r , associate with this pair the set of all those separations in D which distinguish it efficiently and cross as few other separations in D as possible. Let N be the union of all these sets. This set N clearly distinguishes every pair of tangles efficiently, and it is known to be nested [9].

Thus, N is a tree of tangles for G_r . It is canonical in the sense mentioned in [Section 1.1](#), since it is defined purely in terms of invariants of G_r . In particular, it is invariant under the deck transformations of G_r over G (whose group we denote by \mathcal{D}_r). It is thus a well-defined set of separations of G_r that is unique, given G and r . In recognition of the fact that this tree of tangles of G_r distinguishes the lifts of all the local clusters in G and structures them in G_r , we call it, together with \mathcal{D}_r (which reflects this global structure of G_r back to G), the *r -global structure* of G .

This r -global structure of G appears at the level of G as the quotient N/\mathcal{D}_r . This is not a graph. But we show that it can be displayed by a graph, the graph H from [Theorem 1](#), as follows.

Every nested set N of separations of a finite graph can be converted into a tree-decomposition of that graph so that the separations of that graph which correspond to edges of the decomposition tree (the model of the tree-decomposition) are precisely those in N . Infinite graphs in general can have trees of tangles that cannot be represented by a tree-decomposition in this way, because they can have limits under the natural partial ordering of graph separations [12]. For our particular G_r , however, we shall be able to prove that N , our canonical tree of tangles for G_r , does not in fact have problematic limits, and thus gives rise to a canonical tree-decomposition $(T_r, (V_t)_{t \in T_r})$ of G_r . Its orbit space $H = T_r/\mathcal{D}_r$ then is a graph. We

say that H , or more precisely the canonical H -decomposition $(H, (G_h)_{h \in H})$ projected from $(T_r, (V_t)_{t \in T_r})$, displays the r -global structure (N, \mathcal{D}_r) of G . We remark that H is finite if G is finite.

1.4. Overview of the proofs of Theorem 1 and Theorem 2. We will prove the following common generalisation of these two theorems: every connected, locally finite, and quasi-transitive graph G has a graph-decomposition that displays its r -global structure. (A graph is *quasi-transitive* if its automorphism group has finitely many vertex-orbits.) We will see that, if G is quasi-transitive, then G_r is quasi-transitive too.

Our key task will be to prove that the quasi-transitive graph G_r is *tangle-accessible*; this means that all its tangles are pairwise distinguished by separations of some finitely bounded order. For such tangle-accessible G_r our tree of tangles N exists [9], and the separations in N also have finitely bounded order. This implies that N has no problematic limits. In particular, N gives rise to a canonical tree-decomposition of G_r , which will in turn define a canonical graph-decomposition of G that displays its r -global structure.

To show that the quasi-transitive graph G_r is tangle-accessible, we will first use a result of Hamann [22] to see that G_r is *accessible*: that the ends of G_r can be pairwise distinguished by separations of finitely bounded order. By our above considerations there will then exist a canonical tree-decomposition of G_r , of finitely bounded adhesion, that distinguishes all the ends of G_r . We will find a specific such tree-decomposition, \mathcal{T}_{end} , whose infinite parts are 1-ended and themselves quasi-transitive, even just under automorphisms that extend to automorphisms of G_r .

This will allow us to refine \mathcal{T}_{end} in each of its parts so that the arising canonical tree-decomposition \mathcal{T} distinguishes all the tangles of G_r , not just its ends. This \mathcal{T} will again have finitely bounded adhesion. It is not the canonical tree-decomposition of G_r we ultimately seek (the one whose projection to G defines our desired H -decomposition of G), but it witnesses that G_r is tangle-accessible and thus completes our proof.

1.5. A windfall from the proof. The concept of accessibility for locally finite graphs was introduced by Thomassen and Woess [40]. It follows their graph-theoretic characterisation of the finitely generated groups that are accessible, in the sense of Wall [41], by Stallings's group splitting theorem [37]. Hamann, Lehner, Miraftab, and Rühmann [23] extended this characterization to quasi-transitive graphs that need not be Cayley graphs.

In the course of our proof we show that not only the ends of an accessible quasi-transitive graph can be distinguished by separations of bounded order, but all its other tangles can be too:

Theorem 3. *Accessible locally finite, quasi-transitive graphs are in fact tangle-accessible.*

Applied to Cayley graphs of groups, **Theorem 3** yields a structural description of all accessible groups that refines their structural decomposition by Stallings's theorem.

1.6. Structure of the paper. In **Section 2** we collect together notation and basic facts with an emphasis on coverings. In **Section 3** we introduce graph-decompositions, and show how we can obtain them from canonical tree-decompositions via coverings. Local coverings are introduced in **Section 4**, where we also investigate their properties as well as local coverings of some special classes of graphs. In **Section 5** we define formally what it means for a graph-decomposition to 'display the r -global structure' of a graph, and outline what remains to be done to prove **Theorem 1**. The proof is then completed in **Sections 6–8**.

2. TERMINOLOGY

The purpose of this section is to introduce our terminology wherever this may not be obvious or differ from standard usage. Formal definitions of terms not defined here can be found in the extended version of this paper [14], which also contains supplementary proofs and examples.

2.1. Graphs. Our graph-theoretic terminology follows [10], with the following adaptations.

A *graph* G consists of a set $V(G)$ of *vertices* or *nodes*, a set $E(G)$ of *edges* that is disjoint from $V(G)$, and an *incidence* map that assigns to every edge e an unordered pair $uv = \{u, v\}$ of vertices, its *endvertices*. Note that our graphs may have loops and parallel edges: we allow $u = v$, and distinct edges may have the same pair of endvertices.

The triples (e, u, v) and (e, v, u) are the two *orientations* of an edge e ; we think of (e, u, v) as e *oriented from u to v* . When G has no parallel edges, e.g. when G is a tree, we abbreviate (e, u, v) to (u, v) .

A graph is *locally finite* if each of its vertices is incident with only finitely many edges, and *finite* if its vertex and edge sets are both finite. An *isomorphism* between graphs G and H is a pair $\varphi = (\varphi_V, \varphi_E)$ of bijections $\varphi_V: V(G) \rightarrow V(H)$ and $\varphi_E: E(G) \rightarrow E(H)$ that commute with the incidence maps in G and H . We abbreviate both φ_V and φ_E to φ .

2.2. Groups. Our group-theoretic terminology largely follows [31]. Given sets $R \subseteq S$ we write $F(S)$ for the free group that S generates, $\langle R \rangle_{F(S)}^\triangleleft$ for its normal subgroup generated by R , and $\langle S \mid R \rangle$ for the quotient group $F(S)/\langle R \rangle_{F(S)}^\triangleleft$ presented with generators S and relators R . Given a group Γ with a set of generators, we write $\varphi_{\Gamma, S}$ for the canonical homomorphism $F(S) \rightarrow \Gamma$ that extends the identity on S .

2.3. Interplay of graphs and groups. When a group Γ acts on a graph G , we denote the action of an element $g \in \Gamma$ on a vertex or edge $x \in G$ by $x \mapsto g \cdot x$. We write G/Γ for the graph with vertex set $\{\Gamma \cdot v : v \in V(G)\}$ and edge set $\{\Gamma \cdot e : e \in E(G)\}$ in which the edge $\Gamma \cdot e$ has endvertices $\Gamma \cdot u$ and $\Gamma \cdot v$, where u and v are the endvertices of e in G . Note that G/Γ can have loops and parallel edges even if G does not. We call G/Γ the *orbit graph* of the action of Γ on G .

We say that Γ acts *transitively* on G if it does so with only one vertex orbit, and *quasi-transitively* if it does so with only finitely many vertex orbits. The graph G itself is (*quasi-*) *transitive* if its full automorphism group acts (quasi-) transitively on its vertices.

Let Γ be a group and $S \subseteq \Gamma$, usually a set of generators. The *Cayley graph* $\text{Cay}(\Gamma, S)$ is the graph with vertex set Γ , edge set $\{(g, s) : g \in \Gamma, s \in S\}$ and incidence map $(g, s) \mapsto \{g, gs\}$. We allow S to contain inverse pairs of elements. If it contains both s and s^{-1} , say, and $gs = h$ in Γ , then $\text{Cay}(\Gamma, S)$ has two edges incident with g and h : the edge (g, s) and the edge (h, s^{-1}) . Every group acts on its Cayley graphs by left-multiplication, as $h \cdot g := hg$ and $h \cdot (g, s) := (hg, s)$.

2.4. Topological and covering spaces. In our topological notation we follow Hatcher [24]. In particular, when X is a path-connected topological space and $\alpha: [0, 1] \rightarrow X$ a path, we write $\alpha^-: t \mapsto \alpha(1 - t)$ for the path that traverses α backwards. Given two paths α and β with $\alpha(1) = \beta(0)$ we write $\alpha\beta$ for their concatenation, the path which traverses α first and then β , both at double speed. Multiplication in the fundamental group is denoted by $[\alpha][\beta] := [\alpha\beta]$.

For a normal subgroup S of $\pi_1(X, x_0)$ and any $x \in X$ we write S_x for the subgroup $\{[\alpha\gamma\alpha^-] : [\gamma] \in S\}$ of $\pi_1(X, x)$, where α is any path from x to x_0 in X . Since S is normal, this definition is independent of

the choice of α ; in particular, S_x is again normal. When X is a graph G , we call S *canonical* if for every vertex $x \in G$ and for every automorphism φ of G we have $\varphi_*(S_x) = S_{\varphi(x)}$.

When we view a graph G as a 1-complex, as we freely will, we use the term ‘path’ for graph-theoretical paths in the sense of [10], and say *topological path* for continuous maps $[0, 1] \rightarrow G$.

In addition to Hatcher’s [24, Ch.1.3] we use standard covering space terminology as in Jänich [30, Ch.9]. We refer to covering projections $p: (C, \hat{x}_0) \rightarrow (X, x_0)$ as *coverings*, reserving the term *covering space*, or simply *cover*, for C . All covering spaces we consider are connected. The group of deck transformations of a covering p is denoted by $\mathcal{D}(p)$. As generally with automorphisms, we write multiplication in $\mathcal{D}(p)$ as $\varphi \circ \psi =: \varphi\psi$. A normal covering is *canonical* if its characteristic subgroup is canonical.

2.5. Topological paths versus graph-theoretic walks. In this section we provide a graph-theoretical description of the fundamental group of a graph, following Stallings [38], on which our terminology is based.

A *walk* (of length k) in a graph G is a non-empty alternating sequence $v_0e_0v_1 \dots v_{k-1}e_{k-1}v_k$ of vertices and edges in G such that e_i has endvertices v_i and v_{i+1} for all $i < k$. A walk W is *closed* with *base vertex* v_0 if $v_0 = v_k$. We say that W is *based* at v_0 , or simply a closed walk *at* v_0 . A walk is *trivial* if it has length 0. A walk W' is a *subwalk* of W if $W' = v_i e_i v_{i+1} \dots v_{j-1} e_{j-1} v_j$ for some indices i, j with $0 \leq i \leq j \leq k$. The inverse sequence of a walk W is its *reverse*, or W *traversed backwards*; we denote this walk by W^- .

Let G be any connected graph. A walk in G is *reduced* if it contains no subwalk of the form $ueveu$ where u, v are vertices and e is an edge. Note that every walk W can be turned into a reduced walk W' based at the same vertex by iteratively replacing subwalks of the form $ueveu$ with the trivial walk u . We call W' the *reduction* of W and remark that this is well-defined. We call two walks (*combinatorially*) *homotopic* if their reductions are equal. This defines an equivalence relation \sim on the set of all walks in G .

Now let $x_0 \in G$ be any vertex. Let $\mathcal{W}(G, x_0)$ be the set of all closed walks in G that are based at x_0 . Then $\pi'_1(G, x_0) := (\mathcal{W}(G, x_0)/\sim, \cdot)$ is a group, where $[W_1] \cdot [W_2] := [W_1W_2]$ and W_1W_2 is the concatenation of the two walks W_1 and W_2 . Any walk in G defines (up to re-parametrisation) a topological path in the 1-complex G which traverses the vertices and edges of the walk in the same order and direction. This defines a group isomorphism between $\pi'_1(G, x_0)$ and the fundamental group $\pi_1(G, x_0)$ of the 1-complex G [38]. For notational simplicity we will work with combinatorial walks instead of topological paths, and do not always distinguish between walks and the homotopy classes they represent. For example, we may say that a closed walk lies in a subgroup S of $\pi_1(G, x_0)$ when we mean that its combinatorial homotopy class maps to an element of S under the above group isomorphism.

3. GRAPH-DECOMPOSITIONS

In this section we introduce graph-decompositions. We show how graph-decompositions of a graph G can be obtained canonically from canonical tree-decompositions of coverings graphs of G (**Theorem 3.13**).

3.1. Decompositions. Since graph-decompositions generalise tree-decompositions, we recall the notion of tree-decompositions before we introduce the new notion of graph-decompositions.

Definition 3.1 (Tree-decomposition). Let G be a graph, T a tree, and $\mathcal{V} = (V_t)_{t \in T}$ be a family of vertex sets $V_t \subseteq V(G)$ indexed by the nodes t of T . Write G_t for $G[V_t]$, the subgraph of G induced by V_t . The pair (T, \mathcal{V}) is called a *tree-decomposition* of G if it satisfies the following two conditions:

$$(T1) \quad G = \bigcup_{t \in T} G_t;$$

(T2) for every vertex $v \in G$, the graph $T_v := T[\{t \in T : v \in V_t\}]$ is connected.

The vertex sets V_t and the subgraphs $G_t = G[V_t]$ they induce are the *parts* of this tree-decomposition. The tree T is its *decomposition tree*.

In this paper, every part of a tree-decomposition is required to be non-empty. Therefore, if $\mathcal{T} = (T, \mathcal{V})$ is a tree-decomposition of a connected graph G then $V_t \cap V_{t'} \neq \emptyset$ for all $tt' \in E(T)$, by (T2).

A precursor to our notion of graph-decompositions was suggested by Diestel and Kühn [16]. But their notion is not quite the same as ours, and we shall not rely on [16].

Definition 3.2 (Graph-decomposition). Let G and H be graphs, and let $\mathcal{G} = (G_h)_{h \in H}$ be a family of subgraphs $G_h \subseteq G$ indexed by the nodes h of H . The pair (H, \mathcal{G}) is called an *H-decomposition* of G , or more generally a *graph-decomposition*, if it satisfies the following two conditions:

(H1) $G = \bigcup_{h \in H} G_h$;

(H2) for every vertex $v \in G$, the graph $H_v := H[\{h \in H : v \in G_h\}]$ is connected.

The (not necessarily induced) subgraphs G_h are the *parts* of this graph-decomposition. The graph H is its *decomposition graph*. As a short description, we say that H *models* G with parts $(G_h)_{h \in H}$.

If (H, \mathcal{G}) is a graph-decomposition of G such that H is a tree and all the parts G_h are induced subgraphs of G , then (H, \mathcal{G}) defines a tree-decomposition (T, \mathcal{V}) of G with $T := H$ and $V_t := V(G_t)$ for all $t \in T$.

Given an H -decomposition (H, \mathcal{G}) of G as above, we can form its *dual* (G, \mathcal{H}) , where $\mathcal{H} := (H_v)_{v \in G}$. Under mild assumptions, this is a G -decomposition of H . Indeed, let us call (H, \mathcal{G}) *honest* if for every edge hh' of H the intersection $G_h \cap G_{h'}$ is non-empty.

Lemma 3.3 ([14]). *If (H, \mathcal{G}) is an honest graph-decomposition of G into connected parts, then its dual (G, \mathcal{H}) is an honest graph-decomposition of H into connected parts.*

Since the parts G_h of our original graph-decomposition (H, \mathcal{G}) did not have to be induced but the H_v defined in (H2) always are, dualising twice need not return the original decomposition. On the class of graph-decompositions into induced connected subgraphs, however, taking the dual is clearly an involution.

3.2. Properties of graph-decompositions. A *separation* of a set X is an unordered pair $\{A, B\}$ of subsets A, B of X such that $A \cup B = X$. The sets A and B are the *sides* of this separation, their intersection $A \cap B$ is the associated *separator* of X , and $|A \cap B|$ is its *order*. The separation $\{A, B\}$ is *proper* if neither A nor B equals X . The *orientations* of $\{A, B\}$ are the two ordered pairs (A, B) and (B, A) ; these are *oriented separations* of X . We denote the two orientations of a separation s as \vec{s} and \overleftarrow{s} . There are no default orientations: once we have called one of the two orientations \vec{s} , the other will be \overleftarrow{s} , and conversely. A *separation* of a graph G is a separation $\{A, B\}$ of its vertex set $V(G)$ such that G has no edge between $A \setminus B$ and $B \setminus A$.

Let $\mathcal{T} = (T, \mathcal{V})$ be a tree-decomposition of a graph G . Every orientation (t, t') of an edge $e \in T$ induces an oriented separation $\alpha_{\mathcal{T}}(t, t')$ of G , as follows. Let T_t and $T_{t'}$ be the components of $T - e$ containing t and t' , respectively. By (T1), $\alpha_{\mathcal{T}}(t, t') := (A_t, A_{t'})$ is an oriented separation of G with sides $A_t = \bigcup_{s \in T_t} V_s$ and $A_{t'} = \bigcup_{s \in T_{t'}} V_s$. Similarly, $\alpha_{\mathcal{T}}(e) := \{A_t, A_{t'}\}$ is an unoriented separation of G . By (T2), the separator $A_t \cap A_{t'}$ of $\{A_t, A_{t'}\}$ equals $V_t \cap V_{t'}$ [10, Lemma 12.3.1]; this is the *adhesion set* of (T, \mathcal{V}) at the edge $e \in T$. A tree-decomposition has *finitely bounded adhesion* if there exists $K \in \mathbb{N}$ such that all adhesion sets of (T, \mathcal{V}) have order at most K .

Notice that, in the above example, the vertex sets of the components T_t and $T_{t'}$ of $T - e$ formed a separation of the set $V(T)$. It is not a separation of the graph T , since the edge e lies in neither component, but it induces a separation of the graph G . In general, given a graph-decomposition of G modelled on a graph H , separations of the set $V(H)$ induce separations of the graph G in the same way:

Lemma 3.4. *Let $(H, (G_h)_{h \in H})$ be a graph-decomposition of a graph G . Every separation $\{U, W\}$ of the set $V(H)$ induces a separation $\{A, B\}$ of the graph G with $A := \bigcup_{h \in U} V(G_h)$ and $B := \bigcup_{h \in W} V(G_h)$. The separator of $\{A, B\}$ is equal to $(\bigcup_{h \in U \cap W} V(G_h)) \cup \bigcup_{hh' \in F} (V(G_h) \cap V(G_{h'}))$, where $F = E_H(U \setminus W, W \setminus U)$.*

Proof. By (H1), every edge e of G is contained in some part G_h . Then $h \in U$ or $h \in W$, and $e \in G[A]$ or $e \in G[B]$ respectively. In particular, there is no edge in G between $A \setminus B$ and $B \setminus A$. Thus, $\{A, B\}$ is a separation of G .

Let $S := (\bigcup_{h \in U \cap W} V(G_h)) \cup \bigcup_{hh' \in F} (V(G_h) \cap V(G_{h'}))$. It is immediate from the definition of A and B that $A \cap B \supseteq S$. To show $A \cap B \subseteq S$, let $v \in A \cap B$ be any vertex. By definition of A and B there exist $h \in U$ with $v \in G_h$ and $h' \in W$ with $v \in G_{h'}$; in particular, $h, h' \in H_v$. Since H_v is connected (H2), it contains an h - h' path from U to W ; choose h and h' so that this path has minimum length. Since $\{U, W\}$ is a separation of $V(H)$, this means that either $h = h' \in U \cap W$ or $hh' \in F$. In either case, $v \in V(G_h) \cap V(G_{h'}) \subseteq S$ as desired. \square

Recall that an *end* of a graph G is an equivalence class of rays in G , where two rays in G are *equivalent* if for every finite set X of vertices of G the two rays have subrays in the same component of $G - X$. As a first reflection of the relationship between the separations of H and G observed in Lemma 3.4, let us show that the ends of a graph G often correspond naturally to those of its models H .

To define this relationship, we use a tangle-like description of ends known as directions. A *direction* in a graph G is a map f , with domain the set of all finite vertex sets of G , that assigns to every finite vertex set $X \subseteq V(G)$ a component $f(X)$ of $G - X$ so that $f(X) \supseteq f(X')$ whenever $X \subseteq X'$. Every end ω of G defines a direction f_ω in G by letting $f_\omega(X)$ be the component of $G - X$ that contains a subray of one (equivalently: every) ray in ω . The map $\omega \mapsto f_\omega$ is a bijection between the ends of G and its directions [15, Theorem 2.2]. Therefore, in order to define a bijection between the ends of H and G for an H -decomposition of G , it suffices to define a bijection between the directions of H and of G .

Let G be any graph, and let $(H, (G_h)_{h \in H})$ be a graph-decomposition of G with all parts G_h finite. Every direction f in G defines a direction g_f in H , as follows. For every finite set $Y \subseteq V(H)$, the vertex set $X_Y := \bigcup_{h \in Y} V(G_h)$ is finite. By (H2), all of $\bigcup_{v \in f(X_Y)} H_v$ lies in one component of $H - Y$; let $g_f(Y)$ be this component. It is straightforward to verify that g_f is a direction in H .

A graph-decomposition $\mathcal{H} = (H, \mathcal{G})$ of a graph G is *point-finite* if H_v is finite for every vertex v of G .

Lemma 3.5. *For every point-finite and honest H -decomposition \mathcal{H} of a graph G into finite connected parts, the map $f \mapsto g_f$ between the directions of G and those of H is bijective.*

The bijection in Lemma 3.5 between the directions of G and H induces a bijection between their ends, by [15, Theorem 2.2] as mentioned earlier. That bijection is in fact a homeomorphism between the end spaces of G and H ; see [10] for definitions.

Proof of Lemma 3.5. We establish an inverse to the map $f \mapsto g_f$ by constructing a map $g \mapsto f_g$ from the directions of H to those of G such that $f_{g_f} = f$ for all f and $g_{f_g} = g$ for all g .

Let a direction g of H be given. To define f_g , let a finite set $X \subseteq V(G)$ be given. Let Y_X be the set of nodes $h \in H$ whose part G_h contains some vertex of X . Note that Y_X is finite, since \mathcal{H} is point-finite. The subgraph $\bigcup_{h \in g(Y_X)} G_h$ of $G - X$ is connected, because $g(Y_X) \subseteq H - Y$ is connected, the parts $G_h \subseteq G$ are connected, and \mathcal{H} is honest. Let $f_g(X)$ be the unique component of $G - X$ that contains $\bigcup_{h \in g(Y_X)} G_h$.

To check that f_g is a direction, let $X \subseteq X'$ be two finite sets of vertices of G . Then $Y_X \subseteq Y_{X'}$. Since g is a direction, we have $g(Y_X) \supseteq g(Y_{X'})$. Hence $f_g(X) \supseteq \bigcup_{h \in g(Y_X)} G_h \supseteq \bigcup_{h \in g(Y_{X'})} G_h$. Now since $X \subseteq X'$, every component of $G - X'$ is contained in a unique component of $G - X$. As $f_g(X')$ and $f_g(X)$ both contain $\bigcup_{h \in g(Y_{X'})} G_h \neq \emptyset$, the unique component of $G - X$ containing $f_g(X')$ is $f_g(X)$.

For a proof that $f = f_{g_f}$ for every direction f of G , we show that for every finite set $X \subseteq V(G)$ both $f(X)$ and $f_{g_f}(X)$ contain $f(X_{Y_X}) \neq \emptyset$, and must therefore coincide. As $X \subseteq X_{Y_X}$ we have $f(X) \supseteq f(X_{Y_X})$, since f is a direction. For a proof of $f_{g_f}(X) \supseteq f(X_{Y_X})$ let $v \in f(X_{Y_X})$ be given. By (H1) there is an $h \in H$ such that $v \in G_h$. Then $h \in H_v \subseteq g_f(Y_X)$ by definition of g_f and the choice of v , so $G_h \subseteq f_{g_f}(X)$ by definition of f_{g_f} . In particular, $v \in G_h \subseteq f_{g_f}(X)$ as desired.

To show $g = g_{f_g}$ for every direction g of H , we show that for every finite set Y of nodes of H both $g(Y)$ and $g_{f_g}(Y)$ contain $g(Y_{X_Y}) \neq \emptyset$, and must therefore coincide. Since $Y \subseteq Y_{X_Y}$, we have $g(Y) \supseteq g(Y_{X_Y})$ because g is a direction of H . For a proof of $g_{f_g}(Y) \supseteq g(Y_{X_Y})$, let $h \in g(Y_{X_Y})$ be given. Then $V(G_h) \neq \emptyset$ since the parts of \mathcal{H} are connected; pick $v \in V(G_h)$. Now $v \in V(G_h) \subseteq f_g(X_Y)$ by definition of f_g and the choice of h , so $h \in V(H_v) \subseteq g_{f_g}(Y)$ by the definition of g_{f_g} , as desired. \square

Next, let us show that not only do the ends of H and G correspond bijectively, under the assumptions from Lemma 3.5, but the sizes of the separators needed to distinguish these ends correspond too.

A finite-order separation $\{A, B\}$ of G *distinguishes* two ends ω_1, ω_2 of G if for some (and hence every) ray in ω_1 some subray lies in $G[A]$ and for some (and hence every) ray in ω_2 some subray lies in $G[B]$. Following Thomassen and Woess [40], let us call a graph *accessible* if there exists an integer $K \in \mathbb{N}$ such that every two of its ends are distinguished by a separation of order at most K . Compare Section 1.5 for the origin of this notion.

Lemma 3.6. *Let G and H be graphs.*

- (i) *If G has a point-finite and honest H -decomposition into connected parts of finitely bounded size and H is accessible, then so is G .*
- (ii) *If G has an honest H -decomposition into finite connected parts such that the graphs $H_v \subseteq H$ have finitely bounded order and G is accessible, then H too is accessible.*

Proof. For (i) let $n \in \mathbb{N}$ be an upper bound on the size of the decomposition parts, and let $k \in \mathbb{N}$ be such that every two ends of H are distinguished by a separation of order at most k . We show that every two ends of G are distinguished by a separation of order at most nk .

Let ω_1 and ω_2 be two distinct ends of G , and let $f_1 := f_{\omega_1}$ and $f_2 := f_{\omega_2}$ be the two directions of G they define. Consider the directions $g_1 := g_{f_1}$ and $g_2 := g_{f_2}$ of H as defined just before Lemma 3.5. By [15, Theorem 2.2], there exist ends η_1 and η_2 of H which induce these directions g_1 and g_2 . Since H is accessible, it has a separation $\{U_1, U_2\}$ of order at most k such that, for $i = 1, 2$, the end η_i is represented by a ray in U_i , and thus $g_i(U_1 \cap U_2) \subseteq H[U_i]$.

Let us apply Lemma 3.4 to $\{U_1, U_2\}$ to obtain a separation $\{A_1, A_2\}$ of G with $A_i = \bigcup_{h \in U_i} V(G_h)$. Note that $A_1 \cap A_2 = \bigcup_{h \in U_1 \cap U_2} V(G_h)$: as $\{U_1, U_2\}$ is a separation not just of $V(H)$ but of the graph H , the

set F in [Lemma 3.4](#) is empty. Hence $\{A_1, A_2\}$ has order at most nk ; it remains to show that it distinguishes ω_1 and ω_2 .

By definition of f_i , each ω_i has a ray in $f_i(A_1 \cap A_2)$. So it suffices to show that $f_i(A_1 \cap A_2) \subseteq A_i$, for $i = 1, 2$. For every $v \in G$ that is not in $A_1 \cap A_2 = \bigcup_{h \in U_1 \cap U_2} V(G_h)$, the graph H_v avoids $U_1 \cap U_2$, and hence by [\(H2\)](#) lies in a component of $H - (U_1 \cap U_2) \neq \emptyset$. For $v \in f_i(A_1 \cap A_2)$ this component is $g_i(U_1 \cap U_2) \subseteq H[U_i]$, by definition of $g_{f_i} = g_i$. Pick $h \in H_v$ for any such v . Then $v \in G_h$ and $h \in U_i$. This implies $v \in A_i$, as desired, by definition of A_i .

For (ii) let $(G, (H_v)_{v \in G})$ be the honest graph-decomposition of H into connected parts given by [Lemma 3.3](#). Its parts H_v have finitely bounded size by assumption, and it is point-finite since the graphs G_h are finite. Thus, we can apply (i) to $(G, (H_v)_{v \in G})$ to obtain the desired conclusion. \square

3.3. From tree-decompositions to graph-decompositions. In the remainder of this section we show how graph-decompositions of a graph can be obtained from tree-decompositions of its normal covers, and how all these can be chosen canonically. The main result in this section will be [Theorem 3.13](#).

Let $\mathcal{H} = (H, \mathcal{G})$ and $\mathcal{H}' = (H', \mathcal{G}')$ be graph-decompositions of graphs G and G' , respectively. An *isomorphism* between \mathcal{H} and \mathcal{H}' is a pair (φ, ψ) of isomorphisms $\varphi: G \rightarrow G'$ and $\psi: H \rightarrow H'$ such that for every node $h \in H$ we have $\varphi(G_h) = G'_{\psi(h)}$. An isomorphism between \mathcal{H} and itself is an *automorphism* of \mathcal{H} .

A tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of a connected graph G is *regular* if all the separations of G induced by the edges of T are proper: if none has the form $\{A, V(G)\}$. In particular, regular tree-decompositions cannot have empty parts. All tree-decompositions relevant to us will be regular.

We shall be interested particularly in tree-decompositions \mathcal{T} on which the automorphisms φ of the graph being decomposed act naturally, in that they extend to automorphisms (φ, ψ) of \mathcal{T} . These automorphisms ψ are easily seen to be unique when \mathcal{T} is regular:

Lemma 3.7. [\[14\]](#) *If $\mathcal{T} = (T, \mathcal{V})$ is a regular tree-decomposition of a graph G , then for every automorphism φ of G there exists at most one automorphism ψ of \mathcal{T} such that (φ, ψ) is an automorphism of \mathcal{T} .* \square

Let Γ be a group of automorphisms of G . A graph-decomposition $\mathcal{H} = (H, \mathcal{G})$ of G is Γ -*canonical* if Γ acts on \mathcal{H} via $\varphi \mapsto (\varphi, \psi)$ for suitable automorphisms $\psi =: \psi_\varphi$ of H . Once the action of Γ on \mathcal{H} is fixed, or if it is unique by [Lemma 3.7](#), we usually write φ instead of ψ_φ to reduce clutter. If \mathcal{H} is $\text{Aut}(G)$ -canonical, we just call it a *canonical* graph-decomposition of G . When \mathcal{H} is a tree-decomposition, this definition also formalises the notion of canonicity used informally for tree-decompositions in the literature [\[6, 7\]](#).

The graph-decompositions given by the following construction will be the key objects in [Theorem 3.13](#).

Construction 3.8 (Graph-decompositions defined by tree-decompositions via coverings).

Let $p: C \rightarrow G$ be any normal covering of a connected graph G . Let $\mathcal{D} \subseteq \text{Aut}(C)$ be the group of deck transformations, and let $\mathcal{T} = (T, \mathcal{V})$ be any regular \mathcal{D} -canonical tree-decomposition of C . By [Lemma 3.7](#), every $\varphi \in \mathcal{D}$ acts on \mathcal{T} , and in particular on T , so that $\varphi(V_t) = V_{\varphi(t)}$ for all $t \in T$. The map $t \mapsto p(C[V_t])$ is constant on the orbits in $V(T)$ under the action of \mathcal{D} on T , since for every $\varphi \in \mathcal{D}$ we have

$$p(C[V_{\varphi(t)}]) = p(C[\varphi(V_t)]) = p(\varphi(C[V_t])) = p(C[V_t]). \quad (1)$$

Let H be the orbit-graph T/\mathcal{D} of the action of \mathcal{D} on T . Let $\mathcal{G} := (G_h)_{h \in H}$, where $G_h = p(C[V_t])$ for any $t \in h$; this is well-defined by [\(1\)](#). We say that \mathcal{T} *defines* $\mathcal{H} = (H, \mathcal{G})$ via p .

The assumption of regularity for \mathcal{T} in [Construction 3.8](#) is needed only to ensure uniqueness of ψ_φ by [Lemma 3.7](#). This allows us to say that φ , rather than ψ_φ , acts on T , and implies that \mathcal{H} is unique too. The existence of \mathcal{H} as such (without ‘honesty’), as proved in [Lemma 3.9](#), does not require that \mathcal{T} be regular.

Before we show that \mathcal{H} is a graph-decomposition of G , let us note that [Construction 3.8](#) can also be applied to graph-decompositions, rather than tree-decompositions, of the normal cover C . But in this paper we will not need that level of generality.

Lemma 3.9. *Every $\mathcal{H} = (H, \mathcal{G})$ constructed as in [Construction 3.8](#) is an honest graph-decomposition of G .*

Proof. It is straightforward to check [\(H1\)](#).

For a proof of [\(H2\)](#) let v be any vertex of G , and consider any two nodes h_0 and h_1 of H whose parts contain v ; we have to find a walk W from h_0 to h_1 in H with $v \in G_h$ for every node $h \in H$ visited by W . Since v is contained in both G_{h_0} and G_{h_1} , there are nodes t_0 and t_1 of T with $t_i \in h_i$ and vertices $\hat{v}_i \in p^{-1}(v)$ with $\hat{v}_i \in V_{t_i}$ for both $i = 0, 1$. As $p: C \rightarrow G$ is normal, the deck transformations of p act transitively on its fibres, so there is a deck transformation φ of p which maps \hat{v}_0 to \hat{v}_1 . Write $t'_0 := \varphi(t_0)$. By [\(T2\)](#) and since \hat{v}_1 lies in $V_{t'_0} \cap V_{t_1}$, the vertex \hat{v}_1 lies in all parts V_t with $t \in t'_0 T t_1$. As t'_0 and t_1 lie in h_0 and h_1 , respectively, the path $t'_0 T t_1$ in T projects to a walk W from h_0 to h_1 in H such that $v = p(\hat{v}_1)$ is contained in G_h for every node $h \in H$ visited by W .

We have shown so far that $\mathcal{H} = (H, \mathcal{G})$ is a graph-decomposition of G . To check that \mathcal{H} is honest, consider an edge hh' of H . By definition of H there exist adjacent nodes $t \in h$ and $t' \in h'$ in T . As noted in [Section 3.2](#), the set $V_t \cap V_{t'}$ is the separator in a separation $\{A_t, A_{t'}\}$ of C , whose sides are non-empty because \mathcal{T} , being regular, has no empty parts. Since C is connected, as all our covering spaces are, this means that $V_t \cap V_{t'} \neq \emptyset$. Therefore $V(G_h) \cap V(G_{h'}) = p(V_t \cap V_{t'})$ is non-empty too. \square

The graph-decomposition (H, \mathcal{G}) from [Construction 3.8](#) also satisfies the analogue of [\(H2\)](#) for edges:

[\(H3\)](#) For every edge $e \in G$, the graph $H[\{h \in H \mid e \in G_h\}]$ is connected.

The proof is analogous to the proof of [\(H2\)](#) above.

Condition [\(H3\)](#) is a familiar property of all tree-decompositions. For general graph-decompositions not induced by [Construction 3.8](#), however, it can easily fail; see [\[14\]](#) for a simple example.

The graph H obtained in [Construction 3.8](#) can in general be infinite even when G is finite. However this will not happen if \mathcal{T} is point-finite, which our tree-decompositions always will be:

Lemma 3.10. [\[14\]](#) *Graph-decompositions $\mathcal{H} = (H, \mathcal{G})$ constructed as in [Construction 3.8](#) are point-finite if the tree-decomposition \mathcal{T} defining \mathcal{H} is point-finite. If \mathcal{H} is point-finite and G is finite, then H is finite. \square*

For our proof of [Theorem 1](#) we want the graph-decompositions \mathcal{H} of G obtained via [Construction 3.8](#) to be canonical. If all the automorphisms of G are induced via p by automorphisms of C then, as we shall see in [Lemma 3.12](#), the canonicity of the tree-decomposition \mathcal{T} in [Construction 3.8](#) will ensure this. In [Lemma 3.11](#) we provide a sufficient condition on p to ensure this, that all the automorphisms of G ‘lift’ to C in this way.

In general, given a covering $p: C \rightarrow G$ of a connected graph G , a *lift* of an automorphism φ of G is an automorphism $\hat{\varphi}$ of C such that $p \circ \hat{\varphi} = \varphi \circ p$. Lifts of automorphisms are unique up to composition with deck transformations, since for any two lifts $\hat{\varphi}, \hat{\varphi}'$ of φ their composition $\hat{\varphi}' \circ \hat{\varphi}^{-1}$ is a deck transformation.

Lemma 3.11 below says that the automorphisms of G all lift to C if the covering $p: C \rightarrow G$ is canonical (see [Section 2.4](#)):

Lemma 3.11 ([17, Theorem 1]). *Let G be a connected graph, and $p: C \rightarrow G$ a canonical normal covering. Then every automorphism φ of G lifts to an automorphism of C .*

Lemma 3.12. *The graph-decompositions \mathcal{H} from [Construction 3.8](#) are canonical if \mathcal{T} and p are canonical.*

Proof. To prove \mathcal{H} to be canonical, we have to find for every automorphism φ of G an automorphism (φ, ψ) of \mathcal{H} such that $\varphi \mapsto (\varphi, \psi)$ is a group homomorphism.

Let φ be an automorphism of G . As p is canonical, [Lemma 3.11](#) tells us that φ lifts to an automorphism $\hat{\varphi}$ of C . Since \mathcal{T} is canonical, $\hat{\varphi}$ acts on the decomposition tree T , uniquely by [Lemma 3.7](#), as of course does \mathcal{D} . Since $\hat{\varphi}$ is a lift of φ , its action on T is well defined on the orbits of \mathcal{D} in T , i.e., $\psi: \mathcal{D} \cdot t \mapsto \mathcal{D} \cdot \hat{\varphi}(t)$ is a well-defined automorphism of $H = T/\mathcal{D}$. The definition of ψ is also independent of the choice of $\hat{\varphi}$, given φ , since $\hat{\varphi}$ is unique up to composition with deck transformations.

It is routine to check [[14](#)] that our choice of ψ makes $\varphi \mapsto (\varphi, \psi)$ into a group homomorphism from the automorphisms of G to those of \mathcal{H} . \square

Our main result of this section now follows directly from [Lemmas 3.9, 3.10](#) and [3.12](#):

Theorem 3.13. *Let G be any connected graph, and let $p: C \rightarrow G$ be a normal covering of G . Then every regular $\mathcal{D}(p)$ -canonical tree-decomposition \mathcal{T} of C defines a graph-decomposition \mathcal{H} of G via [Construction 3.8](#).*

If \mathcal{T} and p are canonical, then \mathcal{H} is canonical. If \mathcal{T} is point-finite, then so is \mathcal{H} . If \mathcal{H} is point-finite and G is finite, then H is finite. \square

In [[14, 25, 28](#)] we give a number of examples highlighting different aspects of [Construction 3.8](#).

4. LOCAL COVERINGS

Our aim in this section is to construct graph-decompositions $\mathcal{H} = (H, \mathcal{G})$ of a given graph G that display its global structure only, leaving the local aspects of the structure of G to the parts of \mathcal{H} . We shall use the framework of [Construction 3.8](#) to achieve this, but will have to choose the covering $p: C \rightarrow G$ on which it is based specifically to suit our aim. Our choice of this covering will depend only on our intended level of ‘locality’, a parameter r we shall be free to choose in order to set a threshold between ‘local’ and ‘global’.

After defining such ‘local coverings’ formally in [Sections 4.1](#) and [4.2](#), we show in [Section 4.3](#) that these coverings admit a natural hierarchy as their locality parameter r grows.

In [Section 4.5](#) we show that these local covers G_r of finite graphs G always resemble Cayley graphs of finitely presented groups: every G_r has a graph-decomposition into parts of finitely bounded size modelled on a Cayley graph of $\mathcal{D}(p)$, the group of deck transformations of p , which we show is finitely presented.

In [Section 4.6](#) we translate the entire relationship between a graph and its local covers to groups: we show that when G is a locally finite Cayley graph of a finitely generated group Γ , then its local covers G_r are Cayley graphs of finitely presented groups that extend Γ .

4.1. Closed walks which stem from cycles. Let G be any connected graph, and let $x_0 \in G$ be any vertex. We say that a closed walk W in G based at x_0 *stems* from a closed walk Q in G if W can be written as $W = W_0 Q W_0^-$ where the *base walk* W_0 starts at x_0 and ends at the base vertex of Q . A closed walk *once around* a cycle O is a closed walk in O which traverses every edge of O exactly once. If W stems from a closed walk once around a cycle O in G , then W *stems* from O .

Let \mathcal{O} be some fixed set of cycles in G . Given a vertex $x_0 \in G$, denote by $\pi_1^{\mathcal{O}}(G, x_0)$ the subgroup of $\pi_1(G, x_0)$ generated by the closed walks at x_0 that stem from a cycle in \mathcal{O} . For such a walk $W = W_0 Q W_0^-$ its conjugation $U W U^- = (U W_0) Q (U W_0)^-$ with any closed walk U based at x_0 stems from the same cycle as W , as only the base walk has changed. Thus, $\pi_1^{\mathcal{O}}(G, x_0)$ is normal in $\pi_1(G, x_0)$.

We say that a closed walk W in G is *generated* by \mathcal{O} if the following equivalent [14] assertions hold:

- (i) there exists a vertex $x_0 \in G$ such that some walk in $\pi_1^{\mathcal{O}}(G, x_0)$ stems from W ;
- (ii) for every vertex $x_0 \in G$, every closed walk at x_0 in G which stems from W is in $\pi_1^{\mathcal{O}}(G, x_0)$.

In particular, the closed walks at any vertex x_0 that are generated by \mathcal{O} are precisely those in $\pi_1^{\mathcal{O}}(G, x_0)$. Moreover, $S^{\mathcal{O}} := \pi_1^{\mathcal{O}}(G, x_0)$ is canonical as a subgroup of $\pi_1(G, x_0)$ if and only if \mathcal{O} is invariant under the automorphisms φ of G : for every vertex $x \in G$ we have $\varphi_*(S_x^{\mathcal{O}}) = S_{\varphi(x)}^{\varphi(\mathcal{O})} = S_{\varphi(x)}^{\mathcal{O}}$.

We shall need the following well-known relationship between the fundamental group of a graph and the fundamental cycles of its spanning trees (see [24] for a proof). Let T be a spanning tree of a graph G , rooted at a vertex x_0 . Given an edge $e = uv$ not on T , there is a unique cycle in the graph $T + e$, the *fundamental cycle* of e with respect to T . The walk $x_0 T u v T x_0$ and its reverse $x_0 T v u T x_0$ stem from this cycle; they are the only two reduced closed walks at x_0 in $T + e$ that traverse e exactly once. For every such edge e we pick one of these two closed walks and call it the *fundamental closed walk* through e with respect to T .

Lemma 4.1. *Let G be a connected graph, and let T be a spanning tree of G rooted at a vertex x_0 . Then the fundamental group $\pi_1(G, x_0)$ is freely generated by the fundamental closed walks with respect to T . \square*

In particular, every closed walk in G is generated by the set of fundamental cycles with respect to any fixed spanning tree of G . If G is planar, its closed walks are also freely generated by walks around its inner faces [25], which need not arise as the fundamental cycles with respect to any spanning tree. We shall not make use of this fact here.

4.2. Definition of local coverings. Let G be any connected graph, let $x_0 \in G$ be any vertex, and let $r \in \mathbb{N}$. The *r -local subgroup* of $\pi_1(G, x_0)$ is the group

$$\pi_1^r(G, x_0) := \pi_1^{\mathcal{O}_r}(G, x_0)$$

where \mathcal{O}_r denotes the set of all cycles in G of length at most r . When the base point is given implicitly or does not matter, we refer to it simply as $\pi_1^r(G)$.

The *r -local covering* of G , denoted by $p_r: G_r \rightarrow G$, is the unique covering of G with characteristic subgroup $\pi_1^r(G)$. As $\pi_1^r(G)$ is normal and canonical (Section 4.1), so is p_r . We call G_r the *r -local covering graph* of G , or its *r -local cover* for short. In the context of r -local coverings, we may think of r as the *locality parameter*. We call a covering $p: C \rightarrow G$ *local* if p is equivalent to p_r for some $r \in \mathbb{N}$.

If the value of r is clear from the context, we refer to the cycles in \mathcal{O}_r as the *short cycles* in G , and say that a closed walk in G is *generated by short cycles* if it is generated by \mathcal{O}_r . Note that, by standard covering space theory, a walk in G_r is closed if and only if its projection to G is a closed walk in G generated by short cycles.

4.3. Local coverings are ball-preserving. All coverings, by definition, are local homeomorphisms. For r -local coverings of graphs we can extend these local homeomorphisms to larger neighbourhoods. How large exactly depends on r , and our next goal is to quantify this.

Let G be any connected graph, and $\varrho \in \mathbb{N}$. As usual, we write $d_G(x, y)$ for the distance in G between two vertices x and y . Given a vertex $v \in G$, let us write $B_G(v, \varrho/2)$ for the (combinatorial) $(\varrho/2)$ -ball in G around v , the subgraph of G formed by the vertices at distance at most $\varrho/2$ from v and all the edges $xy \in G$ on these vertices that satisfy

$$d_G(v, x) + 1 + d_G(y, v) \leq \varrho.$$

Note that $B_G(v, \varrho/2)$ lacks all the edges xy where x and y have distance exactly $\varrho/2$ from v – which, of course, can occur only when ϱ is even. Instead, it consists of precisely those vertices and edges of G that lie on closed walks at v of length at most ϱ . In particular, $B_G(v, \varrho/2)$ contains all the cycles of length at most ϱ in G that meet v . Given a set X of vertices of G , we write $B_G(X, \varrho/2) := \bigcup_{x \in X} B_G(x, \varrho/2)$ for the $(\varrho/2)$ -ball in G around X .

A covering $p: C \rightarrow G$ is said to *preserve $(\varrho/2)$ -balls* if, for every lift \hat{v} of a vertex $v \in G$, the covering p maps $B_C(\hat{v}, \varrho/2)$ isomorphically to $B_G(v, \varrho/2)$.² In particular, then, p projects all the cycles of length at most ϱ in C isomorphically to cycles of length at most ϱ in G . Note that every covering $C \rightarrow G$ preserves 1-balls.

The preservation of $(\varrho/2)$ -balls can be reduced to local injectivity of the covering map [14]:

Lemma 4.2. *The following two statements are equivalent for coverings $p: C \rightarrow G$ and $\varrho \in \mathbb{N}$:*

- (i) p preserves $(\varrho/2)$ -balls;
- (ii) for every vertex $v \in G$, the distance in C between two distinct lifts of v is greater than ϱ . □

Local coverings, by definition, preserve local structure as encoded in short cycles. Our next lemma says that they preserve local structure also in the more usual sense of copying small neighbourhoods:

Lemma 4.3. *For every $r \in \mathbb{N}$, the r -local covering $p_r: G_r \rightarrow G$ preserves $(r/2)$ -balls.*

Proof. By Lemma 4.2 it is enough to show that no two lifts \hat{v}_1, \hat{v}_2 of a vertex $v \in G$ are connected in G_r by a path of length at most r . The projection of any such path P is a closed walk W at v of length at most r . By a straightforward direct argument, W is generated by closed walks at v stemming from short cycles, and thus lies in $\pi_1^r(G, v)$. Hence W lifts to a closed walk at \hat{v}_1 in G_r , by definition of p_r , which contradicts the fact that its unique such lift P is a path. □

Our r -local covering p_r is universal amongst all the $(r/2)$ -ball-preserving coverings of G :

Lemma 4.4. *For every $(r/2)$ -ball-preserving covering $p: C \rightarrow G$ there is a covering $q: G_r \rightarrow C$ such that $p_r = p \circ q$, and all such coverings q are equivalent.*

Proof. By the Galois correspondence of coverings we only have to show that the characteristic subgroup of p contains that of p_r . The latter is generated by the closed walks W in G that stem from a short cycle O , so it suffices to show that such W lift, via p , to closed walks in C .

Our W has the form $W_0 Q W_0^-$, where Q is a walk once around O . Pick a lift \hat{W} of W to C . This includes, as subwalks, lifts \hat{W}_0, \hat{Q} and \hat{W}_0^- of W_0, Q and W_0^- . Since O is short and p preserves $(r/2)$ -balls, \hat{Q} is closed. Therefore \hat{W}_0^- is the reverse of \hat{W}_0 , and $\hat{W} = \hat{W}_0 \hat{Q} \hat{W}_0^-$ is a closed walk in C . □

²A notion that could be regarded as a precursor of the preservation of $(\varrho/2)$ -balls has been studied by Georgakopoulos [20].

Let us now see how the r -local coverings of G are related for various values of r . When $r \leq r'$, the $(r'/2)$ -balls in G , which the r' -local covering $p_{r'} : G_{r'} \rightarrow G$ preserves by [Lemma 4.3](#), include the $(r/2)$ -balls in G . We may thus apply [Lemma 4.4](#) with $p := p_{r'}$ to obtain a covering $q_{r,r'} : G_r \rightarrow G_{r'}$ such that $p_r = p_{r'} \circ q_{r,r'}$, as shown in [Figure 2](#).




FIGURE 2. The interaction of local coverings for different values of r

Moreover, $G_{r'}$ has its own r -local covering $p_{r,r'} : (G_{r'})_r \rightarrow G_{r'}$. We can compose this with $p_{r'}$ to obtain a covering $p_{r'} \circ p_{r,r'} : (G_{r'})_r \rightarrow G$. This composition of two universal local coverings is itself universal, as it is equivalent to the r -local covering p_r of G :

Lemma 4.5. *Given $r \leq r'$, the coverings $p_{r,r'}$ and $q_{r,r'}$ of $G_{r'}$ are equivalent, while p_r and $p_{r'} \circ p_{r,r'}$ are equivalent as coverings of G . In particular, $(G_{r'})_r$ and G_r are isomorphic graphs.*

Proof. Since both $p_{r'}$ and $p_{r,r'}$ preserve $(r/2)$ -balls, so does their composition $p_{r'} \circ p_{r,r'}$. Choosing this as p in [Lemma 4.4](#) yields a covering $q : G_r \rightarrow (G_{r'})_r$ such that $p_r = p_{r'} \circ p_{r,r'} \circ q$. Since $p_r = p_{r'} \circ q_{r,r'}$ and both p_r and $p_{r'}$ preserve $(r/2)$ -balls, so does $q_{r,r'}$. So by [Lemma 4.4](#) with $p := q_{r,r'}$, we obtain a covering $q' : (G_{r'})_r \rightarrow G_r$ with $p_{r,r'} = q_{r,r'} \circ q'$. Combining all the above, we have

$$p_r = p_{r'} \circ p_{r,r'} \circ q = p_{r'} \circ q_{r,r'} \circ q' \circ q = p_r \circ q' \circ q.$$

As a concatenation of two coverings, $q' \circ q$ is a covering (of G_r by itself), as of course is the identity id_{G_r} . Since $p_r \circ \text{id}_{G_r} = p_r = p_r \circ (q' \circ q)$, as shown above, [Lemma 4.4](#) applied with $p = p_r$ yields that id_{G_r} and $q' \circ q$ are equivalent coverings of G_r . In particular, $q' \circ q$ is a homeomorphism. As q is surjective, being a covering, this implies that q is a homeomorphism, which in turn implies that q' is a homeomorphism. Hence $p_{r,r'}$ and $q_{r,r'}$ are equivalent as coverings of $G_{r'}$ as witnessed by q' , while p_r and $p_{r'} \circ p_{r,r'}$ are equivalent as coverings of G as witnessed by q . \square

Choosing $r = r'$ in [Lemma 4.5](#) shows that taking r -local covers is idempotent: $(G_r)_r = G_r$.

By and large, the connectivity of G_r increases as the locality parameter r grows and G_r changes from a tree, which is minimally connected, to the original graph G . How exactly the connectivity of G_r evolves along the hierarchy of local coverings given by [Lemma 4.5](#) is studied in [\[27\]](#).

4.4. Cycle spaces. As is customary in graph theory, we call the first homology group of a graph its *cycle space*. Taken over the integers it is its *integral cycle space*; taken over $\mathbb{Z}/2\mathbb{Z}$ it is its *binary cycle space*.

To formalise this, let G be any connected graph and $r \in \mathbb{N}$. For every edge e of G pick an arbitrary fixed orientation \vec{e} ; this makes G into a 1-complex. The *integral* or *binary cycle space* of G , then, is the R -module $\mathcal{Z}(G)$ for $R = \mathbb{Z}$ or $R = \mathbb{Z}/2\mathbb{Z}$ that consists of all $E(G) \rightarrow R$ functions with finite support whose sum of the values of the incoming edges at any given vertex equals the sum of the values of its outgoing edges.

Every closed walk $W = v_0 e_0 v_1 \dots v_{k-1} e_{k-1} v_k$ once around a cycle in G induces an element z_W of $\mathcal{Z}(G)$ by letting $z_W(e_i) := 1$ if $\vec{e}_i = (e_i, v_i, v_{i+1})$ and $z_W(e_i) := -1$ if $\vec{e}_i = (e_i, v_{i+1}, v_i)$ and $z_W(e_i) := 0$ elsewhere. Note that the reverse W^- of W induces $z_{W^-} = -z_W$. As is well known and easy to see, the (functions z_W of walks W once around) cycles generate all of $\mathcal{Z}(G)$. We write $\mathcal{Z}_r(G)$ for the submodule of $\mathcal{Z}(G)$ generated by the cycles of length at most r .

Lemma 4.6. *For the r -local cover G_r of G we have $\mathcal{Z}_r(G_r) = \mathcal{Z}(G_r)$.*

Proof. Let us show first that $\pi_1^r(G_r) = \pi_1(G_r)$. By definition of $q_{r,r}$ we have $p_r = p_r \circ q_{r,r}$. By Lemma 4.4 applied with $p = p_r$, it follows from $p_r \circ \text{id}_{G_r} = p_r = p_r \circ q_{r,r}$ that id_{G_r} and $q_{r,r}$ are equivalent coverings of G_r . But $q_{r,r}$ is equivalent also to $p_{r,r}$, by Lemma 4.5. Hence $p_{r,r}$ and id_{G_r} are equivalent coverings of G_r , and thus have the same characteristic subgroups. These are $\pi_1^r(G_r)$ and $\pi_1(G_r)$, respectively.

Since the first homology group of a path-connected space is the abelianisation of its fundamental group, the fact that $\pi_1^r(G_r) = \pi_1(G_r)$ implies $\mathcal{Z}_r(G_r) = \mathcal{Z}(G_r)$. \square

As pointed out for G_r in the last line of the above proof, we have $\mathcal{Z}(G) = \mathcal{Z}_r(G)$ if the fundamental group of G is generated by short cycles. When G is planar the converse holds too; this is shown in [25].

4.5. Local coverings of finite graphs. While local coverings of finite graphs are often infinite, this is not necessarily the case; in [14] we present finite graphs with n -sheeted coverings for any $n \in \mathbb{N}$ constructed by Bowler [2].

Our aim in this section is to prove that the local covers of finite graphs resemble Cayley graphs of finitely presented groups. We first show that every normal cover of a finite graph has a graph-decomposition, into parts of finitely bounded size, modelled on a Cayley graph of its group of deck transformations. This group is always finitely generated. In the case of our local coverings we shall see that it is even finitely presented.

Our first result, the structure theorem for arbitrary normal covers of finite graphs, can be seen as a Švarc-Milnor lemma [33, 39] that provides additional structural information:

Theorem 4.7. *Let $p: C \rightarrow G$ be a normal covering of a finite connected graph G . Then the group \mathcal{D} of all deck transformations of p has a finite generating set S such that C has a graph-decomposition modelled on $D := \text{Cay}(\mathcal{D}, S)$, with connected parts of finitely bounded size. If $\hat{x}_0 \in C$ is a lift of any vertex x of G then these parts C_φ , one for each $\varphi \in \mathcal{D}$, can be chosen as $C_\varphi = B_C(\varphi(\hat{x}_0), |G|)$.*

With these C_φ the graph-decomposition $(D, (C_\varphi)_{\varphi \in \mathcal{D}})$ of C is honest, and the map $f \mapsto g_f$ defined in Section 3.2 induces a bijection between the ends of C and those of D . If \mathcal{D} is an accessible group, then C is an accessible graph.

Proof. Choose $(C_\varphi)_{\varphi \in \mathcal{D}}$ as in the statement of the theorem. As these C_φ depend only on \mathcal{D} , not on the choice of generators for our Cayley graph $\text{Cay}(\mathcal{D}, S)$, we can first verify (H1) and choose S later.

Since C is connected, every vertex is incident with some edge. To verify (H1) it is therefore enough to show that every edge of C is contained in some part C_φ . Every edge of C is a lift $\hat{e} = \hat{u}\hat{v}$ of an edge $e = uv$ of G . These have distance at most $|G| - 1$ from x in G . Pick a shortest u - x path P in G , lift it to a path \hat{P} starting at \hat{u} , and let \hat{x}_1 be the last vertex of \hat{P} . As \hat{P} has length at most $|G| - 1$, this shows that

\hat{e} lies in $B_C(\hat{x}_1, |G|)$. Since p is normal, there exists a deck transformation $\varphi \in \mathcal{D}$ such that $\varphi(\hat{x}_0) = \hat{x}_1$. Thus $\hat{e} \in C_\varphi$, completing the proof of (H1).

Let us now find a finite generating set S of \mathcal{D} such that, for $D = \text{Cay}(\mathcal{D}, S)$, the pair $(D, (C_\varphi)_{\varphi \in \mathcal{D}})$ will satisfy (H2) and thus be a graph-decomposition of C . We shall have to choose S so that, for every lift $\hat{v} \in C$ of a vertex v of G , the subgraph $D_{\hat{v}}$ of D induced by $\{\varphi \in \mathcal{D} : \hat{v} \in C_\varphi\}$ is connected.

To define S , fix of every $v \in G$ one lift \hat{v}_0 . We shall pick enough generators for S to make every $D_{\hat{v}_0}$ connected, and then deduce the same for the other lifts \hat{v} of each v . In fact, we shall add to S for every two nodes of $D_{\hat{v}_0}$ a generator that will make them adjacent in $\text{Cay}(\mathcal{D}, S)$, and hence in $D_{\hat{v}_0}$.

The pairs of nodes of $D_{\hat{v}_0}$ are the pairs $(\varphi, \varphi') \in \mathcal{D}^2$ such that $\hat{v}_0 \in C_\varphi \cap C_{\varphi'}$. To give $D_{\hat{v}_0}$ an edge (φ, s) from φ to φ' we need to add $s = \varphi^{-1}\varphi'$ to S . Let S_v be the set of all these s , setting

$$S_v := \{\varphi^{-1}\varphi' \mid \hat{v}_0 \in C_\varphi \cap C_{\varphi'}\} \subseteq \mathcal{D},$$

and let $S := \bigcup_{v \in V(G)} S_v$. Then the subgraphs $D_{\hat{v}_0}$ of $D = \text{Cay}(\mathcal{D}, S)$ are complete for all $v \in G$.

For a proof that each S_v , and hence S , is finite, it suffices to show that $\hat{v}_0 \in C_\varphi$ for only finitely many φ . For every such φ there is a path in C of length at most $|G|$ from \hat{v}_0 to $\varphi(x_0)$. But $B_C(\hat{v}_0, |G|)$ is finite, and the $\varphi(x_0)$ are distinct for different φ , by the uniqueness of path lifting. So there are only finitely many such φ , as required.

To prove (H2) for lifts \hat{v} other than \hat{v}_0 , recall from Section 2.3 that \mathcal{D} acts on D by left-multiplication. Since our covering p is normal, there exists $\psi \in \mathcal{D}$ such that $\hat{v} = \psi(\hat{v}_0)$. Then $\psi \cdot D_{\hat{v}_0} = \{\psi\varphi \mid \varphi \in D_{\hat{v}_0}\}$, and any two of these nodes $\psi\varphi$ and $\psi\varphi'$ are joined by the edge $(\psi\varphi, s)$ for $s = \varphi^{-1}\varphi' \in S_v$. So it suffices to check that the node set of $\psi \cdot D_{\hat{v}_0}$ equals that of $D_{\psi(\hat{v}_0)} = D_{\hat{v}}$.

By definition, $D_{\psi(\hat{v}_0)}$ consists of those $\chi \in \mathcal{D}$ for which $\psi(\hat{v}_0) \in C_\chi$, i.e., for which $\psi(\hat{v}_0)$ has distance at most $|G|$ from $\chi(\hat{x}_0)$. With $\chi =: \psi\varphi$ the latter is equivalent to \hat{v}_0 having distance at most $|G|$ from $\varphi(\hat{x}_0)$, i.e. to $\varphi \in D_{\hat{v}_0}$. Hence

$$\chi \in D_{\psi(\hat{v}_0)} \Leftrightarrow \varphi \in D_{\hat{v}_0} \Leftrightarrow \psi\varphi \in \psi \cdot D_{\hat{v}_0} \Leftrightarrow \chi \in \psi \cdot D_{\hat{v}_0}$$

as desired. This completes the proof that $(D, (C_\varphi)_{\varphi \in \mathcal{D}})$ is a graph-decomposition.

For a proof that S generates \mathcal{D} , we have to show that $D = \text{Cay}(\mathcal{D}, S)$ is connected. This follows from the connectedness of C by Lemma 3.4, since the parts C_φ of our graph-decomposition are non-empty.

Next, we check that our graph-decomposition $(D, (C_\varphi)_{\varphi \in \mathcal{D}})$ is honest: that $C_\varphi \cap C_{\varphi'} \neq \emptyset$ for any adjacent nodes φ, φ' of D . As $D = \text{Cay}(\mathcal{D}, S)$ there exists $s \in S$ such that $\varphi' = \varphi s$. This s lies in some S_v . This means, by definition of S_v , that there exist $\varphi_0, \varphi'_0 \in \mathcal{D}$ such that $s = \varphi_0^{-1}\varphi'_0$ and $\hat{v}_0 \in C_{\varphi_0} \cap C_{\varphi'_0}$. Let $\psi := \varphi\varphi_0^{-1}$. Then $\varphi = \psi\varphi_0$ and $\varphi' = \varphi s = \psi\varphi_0 s = \psi\varphi'_0$. Hence $\hat{v}_0 \in C_{\varphi_0} \cap C_{\varphi'_0}$ implies that $\psi(\hat{v}_0) \in C_{\psi\varphi_0} \cap C_{\psi\varphi'_0} = C_\varphi \cap C_{\varphi'}$, since the action of \mathcal{D} on itself by left-multiplication commutes with its natural action on C . In particular, $C_\varphi \cap C_{\varphi'} \neq \emptyset$, as required for our honesty proof.

We have shown that our graph-decomposition $(D, (C_\varphi)_{\varphi \in \mathcal{D}})$ is honest and has finite connected parts. In our proof that the sets S_v are finite we saw that each vertex of C lies in C_φ for only finitely many $\varphi \in \mathcal{D}$: that $(D, (C_\varphi)_{\varphi \in \mathcal{D}})$ is point-finite. By Lemma 3.5, therefore, the directions of C and of $D = \text{Cay}(\mathcal{D}, S)$ are in bijective correspondence, and hence so are their ends [15, Theorem 2.2]. The ends of \mathcal{D} , by definition, are those of D .

Finally, since C covers the finite graph G it has finitely bounded degrees. The parts $C_\varphi = B_C(\varphi(\hat{x}_0), |G|)$ thus have finitely bounded size. If the group \mathcal{D} is accessible then so is its Cayley graph D [40], and hence by Lemma 3.6 so is C . \square

The groups of deck transformations of our local coverings of finite graphs are not only finitely generated but even finitely presented:

Proposition 4.8. *Let G be a finite graph, and let $r \in \mathbb{N}$. Then the group \mathcal{D}_r of deck transformations of the r -local covering of G is finitely presented.*

Proof. Fix a vertex $x_0 \in G$. The fundamental group $\pi_1(G, x_0)$ of G is finitely presented by [Lemma 4.1](#). To show that $\mathcal{D}_r = \pi_1(G, x_0)/\pi_1^r(G, x_0)$ is finitely presented, it remains to prove that $\pi_1^r(G, x_0)$ is the normal closure of a finite set of closed walks in G .

Recall that $\pi_1^r(G, x_0)$ is generated by the closed walks at x_0 stemming from short cycles in G . There are only finitely many such cycles, but infinitely many such walks. However, the walks stemming from the same short cycle O are all conjugates of each other in $\pi_1(G, x_0)$. Hence we may pick one of them for each O , finitely many in total, and their normal closure in $\pi_1(G, x_0)$ will be exactly $\pi_1^r(G, x_0)$. \square

Corollary 4.9. *Local covers of finite graphs are accessible. They have graph-decompositions into connected parts of finitely bounded size modelled on Cayley graphs of finitely presented groups.*

Proof. Finitely presented groups are accessible [\[18\]](#). Apply [Theorem 4.7](#) and [Proposition 4.8](#). \square

We remark that arbitrary normal covers of finite graphs can be inaccessible [\[14\]](#).

4.6. Local coverings of Cayley graphs. When our graph G is itself a Cayley graph, of some finitely generated (but not necessarily finite) group Γ , say, we can say even more about its local covers than we could in [Section 4.5](#). These covers are now Cayley graphs too, of groups extending Γ :

Theorem 4.10. *Let G be a connected locally finite Cayley graph of some group Γ , and $r \in \mathbb{N}$. Then G_r is a connected locally finite Cayley graph of a finitely presented group Γ_r of which Γ is a quotient.*

For our proof of [Theorem 4.10](#) we need a tool from combinatorial group theory.

Let G be a Cayley graph of a group Γ with respect to a finite generating set S . The fundamental group $\pi_1(G, x_0)$ of G is then isomorphic to the kernel of the canonical homomorphism $\varphi_{\Gamma, S}: F(S) \rightarrow \Gamma$, where $F(S)$ is the free group generated by S . Let us define such an isomorphism explicitly.

The *edge-labelling* $\ell: E(G) \rightarrow S$ of a Cayley graph $G = \text{Cay}(\Gamma, S)$ is defined as $\ell(e) := s$ when $e = (g, s)$. We extend ℓ to a map from the walks in G to $F(S)$ by mapping a walk $W = g_0 e_0 g_1 \dots g_{k-1} e_{k-1} g_k$ to the group element $\ell(e_0)^{\epsilon_0} \dots \ell(e_{k-1})^{\epsilon_{k-1}}$ of $F(S)$, where ϵ_i is 1 if $e_i = (g_i, s)$ for some $s \in S$ and -1 otherwise. Since two walks map to the same element of $F(S)$ under ℓ if their reductions coincide, ℓ defines a group homomorphism $\ell_*: \pi_1(G, x_0) \rightarrow F(S)$. This ℓ_* is our explicit isomorphism between $\pi_1(G, x_0)$ and $\ker(\varphi_{\Gamma, S}) \leq F(S)$:

Lemma 4.11 ([\[32, Theorem 1.6\]](#)). *Let G be a Cayley graph of a group Γ with respect to a finite generating set S , and let ℓ be its edge-labelling. Fix any vertex x_0 of G . Then ℓ_* is an isomorphism from the fundamental group $\pi_1(G, x_0)$ to the kernel of the canonical homomorphism $\varphi_{\Gamma, S}: F(S) \rightarrow \Gamma$.*

Proof of [Theorem 4.10](#). Let G be the Cayley graph of Γ with respect to the finite generating set S . Then Γ has the group presentation $\langle S \mid R \rangle$ with $R = \ker(\varphi_{\Gamma, S})$. We aim to show that G_r is the Cayley graph of a group Γ' with the same set S of generators as G , but with a different and finite set $R' \subseteq R$ of relators. We shall first construct Γ' and R' , then check that $\text{Cay}(\Gamma', S)$ is a covering graph of G , and finally show that this covering is equivalent to our r -local covering of G .

In $G = \text{Cay}(\Gamma, S)$ let us fix the identity $1 \in \Gamma$ as our basepoint x_0 . Since G is locally finite, there are only finitely many cycles of length at most r that contain x_0 . Hence the set of closed walks at x_0 that only go once around such a cycle is finite too; let $R' \subseteq F(S)$ be its image under ℓ . By [Lemma 4.11](#) we have $R' \subseteq \ell_*(\pi_1(G, x_0)) = \ker(\varphi_{\Gamma, S}) \subseteq F(S)$. As R' is finite, the group $\Gamma' := \langle S \mid R' \rangle$ is finitely presented.

Let us construct a covering $p': G' \rightarrow G$ for $G' = \text{Cay}(\Gamma', S)$. Since $R' \subseteq \ker(\varphi_{\Gamma, S})$, we have $\ker(\varphi_{\Gamma', S}) = \langle R' \rangle_{F(S)}^{\triangleleft} \subseteq \ker(\varphi_{\Gamma, S})$. Hence $\Gamma = F(S)/\ker(\varphi_{\Gamma, S})$ is a quotient of $\Gamma' = F(S)/\ker(\varphi_{\Gamma', S})$. The quotient map $p': \Gamma' \rightarrow \Gamma$ between these groups, the vertex sets of their Cayley graphs G' and G , extends to a covering $p': G' \rightarrow G$ by $(g', s) \mapsto (p'(g'), s)$ for edges, since both Cayley graphs use the same set S as edge labels. Explicitly, we have $\ell \circ p' = \ell'$ for the edge-labelling ℓ' of G' .

Let us show that p' is equivalent to our r -local covering $p_r: G_r \rightarrow G$. We do so by showing that their characteristic subgroups coincide. By [Lemma 4.11](#), this will be the case if $\ell'_*(\pi_1(G', \hat{x}_0))$ and $\ell_*(\pi_1^r(G, x_0))$ are the same subgroup of $F(S)$, where \hat{x}_0 is any element in $p'^{-1}(x_0)$. Let R_r be the image under ℓ of the set of closed walks at x_0 which stem from a cycle in G of length at most r . Note that R_r is usually infinite, whereas $R' \subseteq R_r$ is finite. Again by [Lemma 4.11](#) we have $\langle R' \rangle_{F(S)}^{\triangleleft} = \ell'_*(\pi_1(G', \hat{x}_0))$ and $\langle R_r \rangle_{F(S)}^{\triangleleft} = \ell_*(\pi_1^r(G, x_0))$. It thus remains to show that $\langle R_r \rangle_{F(S)}^{\triangleleft} = \langle R' \rangle_{F(S)}^{\triangleleft}$.

Since $R' \subseteq R_r$, we immediately have $\langle R' \rangle_{F(S)}^{\triangleleft} \subseteq \langle R_r \rangle_{F(S)}^{\triangleleft}$. For the converse inclusion let $t \in R_r$ be given. Then there exists a reduced closed walk W at x_0 which stems from a cycle O of length at most r such that $\ell(W) = t$. This W consists of a base walk W_0 and a closed walk Q once around O , so that $W = W_0 Q W_0^-$. Let $g \in V(G) = \Gamma$ be the base vertex of Q . Left-multiplication by g^{-1} defines an automorphism $\psi: h \mapsto g^{-1}h$ of G . Now $\psi(Q)$ is a closed walk, based at $\psi(g) = 1 = x_0$, once around the cycle $\psi(O)$ of length at most r . Since ψ is an automorphism of a Cayley graph, the labellings on the edges of O and $\psi(O)$ agree, so that $\ell(Q) = \ell(\psi(Q)) \in R'$. Thus, $t = \ell(W) = \ell(W_0)\ell(Q)\ell(W_0^-) = \ell(W_0)\ell(Q)\ell(W_0)^{-1}$ is contained in the normal subgroup $\langle R' \rangle_{F(S)}^{\triangleleft}$. This yields $R_r \subseteq \langle R' \rangle_{F(S)}^{\triangleleft}$ and completes the proof. \square

Corollary 4.12. *Local covers of connected locally finite Cayley graphs are accessible.*

Proof. Finitely presented groups are accessible [18], and hence so are their locally finite Cayley graphs [40]. Apply [Theorem 4.10](#). \square

5. DISPLAYING GLOBAL STRUCTURE

In this section we make the statements of [Theorems 1](#) and [2](#) precise and give an overview of their proof. This will include a proof of [Theorem 3](#).

In [Section 5.1](#) we recall how a tree of tangles can describe the tree-like aspects of the global structure of any graph. In [Section 5.2](#) we apply this to the typically infinite graphs G_r that reflect the r -local structure of the finite graphs G given in [Theorem 1](#), to describe their (tree-like) r -global structure. Projected back to G , this r -global structure of G_r ceases to be infinite and tree-like, but defines what we call the (finite) r -global structure of G .

[Section 5.2](#) ends with a precise definition of the term of ‘ r -global structure’ used in [Theorems 1](#) and [2](#), and of what it means for another graph to ‘display’ that structure. We proceed from there to give an outline of the proof of [Theorem 1](#). This includes a common generalisation, [Theorem 5.5](#), of [Theorems 1](#) and [2](#). The outline will serve as our road map when we prove [Theorems 1](#) and [2](#) (by way of proving [Theorem 5.5](#)) in [Sections 6–8](#).

5.1. Separations, tangles, and trees of tangles. In this section we recall basic facts and terminology about tangles; for a more detailed account see [10, Ch.12.5].

Let G be a graph. Given a set S of separations of G , we write \vec{S} for the set of orientations of separations in S , two for every $s \in S$. An *orientation* O of S contains for every separation in S exactly one of its two orientations. Recall that $(A, B) \leq (C, D)$ for two oriented separations (A, B) and (C, D) of G if $A \subseteq C$ and $B \supseteq D$. If O does not contain an oriented separation (B, A) of G whenever $(A, B) \leq (C, D) \in O$, then O is *consistent*. A set N of separations of G is *nested* if every two of its elements, $\{A, B\}$ and $\{C, D\}$ say, have orientations $(A, B) \leq (C, D)$. Separations that are not nested are said to *cross*.

Let Γ be a group acting on a graph G . A set S of separations of G is Γ -*canonical* if for every $g \in \Gamma$ and every $\{A, B\} \in S$ we have $g \cdot \{A, B\} := \{g \cdot A, g \cdot B\} \in S$. If S is $\text{Aut}(G)$ -canonical, we call it *canonical*.

Let us now define tangles. Given any $k \in \mathbb{N} \cup \{\aleph_0\}$, let S_k denote the set of all separations of G of order $< k$. A k -*tangle* in G is an orientation τ of S_k such that for all $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \tau$, not necessarily distinct, we have $G[A_1] \cup G[A_2] \cup G[A_3] \neq G$. A *tangle* is a k -tangle for some k . Note that every tangle is consistent.

Examples of tangles include blocks and ends. A k -*block* in G is a maximal set of at least k vertices that is included in one of the two sides of each separation in S_k . It *orients* that separation *towards* this side: as (A, B) if it lies in B . Every k -block X thus defines an orientation τ_X of S_k . If $|X| > 3(k-1)/2$, then τ_X is a tangle [5, Theorem 6.1]. By the maximality in the definition of a k -block, distinct k -blocks define different tangles.

Given an end ω of G , let f_ω be the corresponding direction of G (see Section 3.2). Orienting every separation $\{A, B\} \in S_{\aleph_0}$ towards the side that contains $f_\omega(A \cap B)$, we obtain a tangle τ_ω of S_{\aleph_0} [11]. As in the case of blocks, distinct ends define different tangles.

A separation of G *distinguishes* two tangles in G if they orient it differently. The separation distinguishes these two tangles *efficiently* if no separation of smaller order distinguishes them. Two tangles are *distinguishable* if some separation distinguishes them; note that any two k -tangles for the same k are distinguishable. A set N of separations *distinguishes* a set of tangles (efficiently) if every two distinguishable tangles in this set are (efficiently) distinguished by some separation in N . A *tree of tangles* of G , in this paper, is a nested set of separations of G which efficiently distinguishes the entire set of tangles in G .

A separation of G is *relevant* if it efficiently distinguishes some two tangles τ and τ' in G . One can show that such separations $\{A, B\}$ are *tight*: for $X := A \cap B$ the graph $G - X$ has components $C \subseteq G[A]$ and $D \subseteq G[B]$ such that every vertex in X sends edges to both C and D . Indeed, these are the unique such components for which $(V(G-C), V(C) \cup X) \in \tau$ and $(V(G-D), V(D) \cup X) \in \tau'$ or vice versa [10, Ex. 43(iii)].

Note that a finite-order separation of G distinguishes the tangles induced by two ends of G if and only if it distinguishes these ends in the usual sense that they contain rays on different sides of that separation. It distinguishes these two ends *efficiently* if no separation of lower order distinguishes them.

Let us remember all the above properties of trees of tangles for later use in Section 8:

Lemma 5.1. *Let N be a tree of tangles of G in which every separation is relevant. Then N efficiently distinguishes every two ends of G . Every separation in N has finite order and is tight.* \square

5.2. The r -global structure of G . In general, a locally finite graph G may have several canonical trees of tangles. However, there is one specific canonical tree of tangles $N(G)$ that is particularly natural, which

we therefore fix as an invariant of those locally finite graphs in which it exists. In this section we describe the construction of $N(G)$, and then show in which sense it is particularly natural.

Recall [23, 40] that a graph G is *accessible* if there exists a number $K \in \mathbb{N}$ such that every two ends of G are distinguished by a separation of order at most K . Extending this notion from ends to tangles, we call a graph G *tangle-accessible* if there exists a number $K \in \mathbb{N}$ such that every two distinguishable tangles in G are distinguished by a separation of order at most K .

Let G be a connected, locally finite, and tangle-accessible graph. Denote by D the set of all finite-order separations of G that distinguish some two tangles in G efficiently. Since G is tangle-accessible, all the separations in D have order at most some fixed $K \in \mathbb{N}$, and hence they each cross only finitely many other separations in D [19, proof of Proposition 6.2].

Definition 5.2. For every pair τ, τ' of distinguishable tangles in G let $N_{\tau, \tau'}$ be the set of all separations that efficiently distinguish τ from τ' and, among these, cross as few separations in D as possible. Let $N(G)$ denote the union of all these sets $N_{\tau, \tau'}$.

Note that the separations in $N_{\tau, \tau'}$ in Definition 5.2 are relevant. The set $N(G)$ is clearly canonical, it distinguishes all the tangles in G , and its elements have finitely bounded order. It is also still nested:

Lemma 5.3 ([9, Theorem 4.2]). *For every connected, locally finite, and tangle-accessible graph G , the set $N(G)$ of separations defined in Definition 5.2 is a canonical tree of tangles. All the separations in $N(G)$ are relevant; in particular, they have finitely bounded order.*

We sometimes refer to $N(G)$ informally as *the* canonical tree of tangles of G .

We shall prove in Section 7 that the r -local covers G_r of the graphs G we consider in Theorems 1 and 2 satisfy the premise of Lemma 5.3. We shall see that they have unique canonical tree-decompositions $\mathcal{T} = (T, \mathcal{V})$ such that $\alpha_{\mathcal{T}}: E(T) \rightarrow N(G_r)$ is a bijection. We say that $N(G_r)$ *induces* this tree-decomposition \mathcal{T} of G_r .

To see in which sense $N(G)$ is particularly natural, consider any tree of tangles N in G . By definition, N has to contain for every pair τ, τ' of tangles a separation $s(\tau, \tau')$ that distinguishes them efficiently. If we want N to be canonical, however, we cannot simply pick for N just any such separation $s(\tau, \tau')$ for each pair τ, τ' of tangles: our choice has to be definable in terms of invariants of G . This may force us to pick not just one such separation $s(\tau, \tau')$ for each pair τ, τ' , but a set $N(\tau, \tau')$ of more than one separation. As $N(\tau, \tau')$ is to be included in N , it will be a nested subset of the set $D(\tau, \tau')$ of separations that distinguish τ from τ' efficiently. Our set $N(G)$ now chooses these $N(\tau, \tau')$ with reference to $D(\tau, \tau')$ alone: they consist of *all* the elements of $D(\tau, \tau')$ that are the most promising candidates for inclusion in $N(\tau, \tau')$, in that they cross as few other elements of $D(\tau, \tau')$ as possible.

As described in Section 4, the r -local cover G_r of a graph G agrees with G on its r -local structure, the structure encoded in the r -local subgroup $\pi_1^r(G)$ of $\pi_1(G)$. The remaining structure of G is unfolded in G_r in a tree-like way, which encodes it in conjunction with the group $\mathcal{D}_r = \pi_1(G)/\pi_1^r(G)$ of deck transformations of G_r over G . This tree-like global structure of G_r is captured by its tree of tangles $N(G_r)$. We refer to $N(G_r)$ and \mathcal{D}_r together as the r -global structure of G :

Definition 5.4. The r -global structure of a graph G whose r -local cover G_r is tangle-accessible is the pair $(N(G_r), \mathcal{D}_r)$, where \mathcal{D}_r is the group of deck transformations of $p_r: G_r \rightarrow G$ and $N(G_r)$ is our canonical tree of tangles of G_r from Definition 5.2.

The r -global structure of a group Γ presented with a finite set S of generators is the r -global structure of its Cayley graph $\text{Cay}(\Gamma, S)$.

The unique canonical graph-decomposition of G that *displays* its r -global structure $(N(G_r), \mathcal{D}_r)$ is that defined via p_r by the canonical tree-decomposition which $N(G_r)$ induces on G_r , as in [Construction 3.8](#). \square

[Definition 5.4](#) rests on the assumption that G_r is tangle-accessible. We shall prove that this is the case whenever G is finite, as in [Theorem 1](#), or a Cayley graph as in [Theorem 2](#).

Note that when a graph-decomposition $\mathcal{H} = (H, \mathcal{G})$ displays the r -global structure of a graph G , then H comes with an edge labelling that reveals additional internal structure of H , which in turn displays additional connectivity aspects of G not reflected in the graph structure of H itself.

Indeed, recall that every edge e of the tree T from the tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ that defines \mathcal{H} as in [Construction 3.8](#) maps to a separation $\alpha(e)$ of G_r . This separation has some order k_e . When one analyses trees of tangles in graph minor theory, one often contracts in T all those edges e whose label k_e is at least some chosen threshold k . This results in a tree T_k that is a minor of T , and whose edges e map to separations $\alpha(e)$ of order $< k$. These still distinguish all the k -tangles of the graph decomposed, in our case G_r . These trees T_k form a chain of minors of T as k grows: $T_1 \preceq T_2 \preceq \dots \preceq T$.

This structure is now inherited by $H = T/\mathcal{D}$. Its edges arise from (\mathcal{D} -orbits of) edges e of T , which carry labels k_e that specify the orders of separations $\alpha(e)$ of G_r . We may interpret these as local separations of G : they tell us how the parts G_h of \mathcal{H} are separated locally from each other in G . (In our original example from the Introduction, these local separators are the 2-sets in which two adjacent K^5 s overlap.) And as with those tree minors T_k of T , contracting in H the edges with labels at least some chosen threshold k gives us a minor H_k of H . We thus obtain a hierarchy $H_1 \preceq H_2 \dots$ of minors of H that defines an internal structure of H , which in turn displays more details of the connectivity structure of G than H does alone. This internal structure is explored further in [\[27, 28\]](#).

It is sometimes desirable that the parts G_h of the graph-decomposition that displays the r -global structure of G should inherit more properties of the parts $G_r[V_t]$ of \mathcal{T} whose projections they are than the projection map p_r can transmit. Such additional information is encoded in what we call a *blueprint* of G : a family whose members are not themselves subgraphs of G but which map homomorphically to the parts G_h . We explore the notion of blueprints and the properties of the parts of \mathcal{T} that they preserve in [\[29\]](#). This yields, among other things, a new proof of Carmesin's local 2-separator theorem [\[4\]](#) via local coverings.

While the r -global structure of a finite graph G is displayed by a graph-decomposition (H, \mathcal{G}) of G where H is another finite graph, this decomposition of G is obtained indirectly via its usually infinite r -local covering. One can ask, then, whether it is possible to characterise (H, \mathcal{G}) directly in terms of G . As a proof of concept for this type of problem, we obtain in [\[26\]](#) such a direct characterisation of the r -global structure of the locally finite graphs that are ' r -locally chordal', those whose parts G_h as above are complete.

5.3. Road map for the proof of main result. The remainder of this paper is devoted to the proof of [Theorem 1](#), for which we now give an outline. We shall prove [Theorem 1](#) by way of proving a stronger result, [Theorem 5.5](#), which will include [Theorem 2](#) at no extra cost. A proof of [Theorem 3](#) will also be included on the way.

As we will see in [Section 7.1](#), every nested set of separations of a graph, all of finitely bounded order, induces a tree-decomposition of that graph. As seen in [Construction 3.8](#) and [Theorem 3.13](#), canonical tree-decompositions of local covers define canonical graph-decompositions of the graph they cover. So once

we know that our canonical set $N(G_r)$ from [Definition 5.2](#) is defined, for which we only need that G_r is tangle-accessible ([Lemma 5.3](#)), it will induce the tree-decomposition of G_r that defines the canonical graph-decomposition of G which displays its r -global structure.

Let us summarise these observations as the point of departure for the proofs of all our main results:

Reduction 1. *It suffices to show that G_r is tangle-accessible.*

In order to show that G_r is tangle-accessible we study another canonical tree of tangles of G_r , one whose existence was proved by Elbracht, Kneip and Teegen [[19](#)] ([Lemma 8.4](#)). In contrast to $N(G_r)$ this tree of tangles $N'(G_r)$ is defined for all locally finite graphs, without any accessibility requirements. Correspondingly, the separations in $N'(G_r)$ need not, prima facie, have finitely bounded order. However if they do, then G_r is tangle-accessible. Hence:

Reduction 2. *It suffices to show that the separations in $N'(G_r)$ have finitely bounded order.*

Instead of trying to bound the order of the separations in $N'(G_r)$ directly, we will prove that $N'(G_r)$ exhibits a certain structural property: the property of being *exhaustive*. This means that for every infinite sequence $(A_0, B_0) < (A_1, B_1) < \dots$ in $\vec{N}'(G_r)$ we have $\bigcap_{i \in \mathbb{N}} B_i = \emptyset$. Sets of tight separations of finitely bounded order in locally finite graphs are easily seen to be exhaustive. The converse fails in general. However we will show that if $N'(G_r)$ is exhaustive, then \mathcal{D}_r acts on $N'(G_r)$ with only finitely many orbits. In particular, then, the separations in $N'(G_r)$ will have finitely bounded order. Thus:

Reduction 3. *It suffices to show that $N'(G_r)$ is exhaustive.*

We will prove that $N'(G_r)$ is exhaustive in two steps, as follows.

In the first step we show that G_r is accessible in the traditional sense: that the separations in $N'(G_r)$ that efficiently distinguish two ends of G_r have finitely bounded order. We denote this set as $N'_{\text{end}} \subseteq N'(G_r)$. The fact that the separations in N'_{end} have finitely bounded order, and hence are exhaustive because they are clearly tight, will be the key ingredient of the remainder of our proof.

In the second step we show that $N'(G_r)$ is exhaustive. For this let $(A_0, B_0) < (A_1, B_1) < \dots$ be any infinite sequence in $\vec{N}'(G_r)$; we have to show $\bigcap_{i \in \mathbb{N}} B_i = \emptyset$. Since N'_{end} is exhaustive, it suffices to find in \vec{N}'_{end} an infinite sequence $(C_0, D_0) < (C_1, D_1) < \dots$ which our earlier sequence dominates: one such that for every $i \in \mathbb{N}$ there exists $j \in \mathbb{N}$ with $(C_i, D_i) \leq (A_j, B_j)$. If there is no such sequence we find a part V'_t of the tree-decomposition of G_r induced by N'_{end} such that $A_j \cap B_j \subseteq V'_t$ for all large enough j . We then enlarge V'_t to a certain set V_t such that, for every j_0 large enough, the subsequence $(A_j, B_j)_{j \geq j_0}$ induces an infinite sequence of tight separations of $G[V_t]$. Moreover, $G[V_t]$ will have the following properties.

The graph $G[V_t]$ will be connected; it will have exactly one end; and the stabiliser of t in \mathcal{D}_r , the group of the deck transformations of G_r , will act quasi-transitively on $G[V_t]$. These properties combined will imply that $G[V_t]$ has no infinite sequence of separations such as that induced by our $(A_j, B_j)_{j \geq j_0}$, a contradiction. Hence $N'(G_r)$ will be exhaustive, and our proof will be complete.

Following this road map we will be able to prove [Theorems 1](#) and [2](#) together, by proving their following common strengthening:

Theorem 5.5. *Let G be a connected, locally finite, and quasi-transitive graph, and let $r > 0$ be an integer. Then G has a unique canonical and point-finite graph-decomposition which displays its r -global structure.*

Note that finite graphs, such as the G in [Theorem 1](#), are trivially quasi-transitive, and so are locally finite Cayley graphs such as the $\text{Cay}(\Gamma, S)$ in [Theorem 2](#).

Our proof of [Theorem 5.5](#) will implement the strategy outlined above by reversing the order of its steps, as follows. We start in [Section 6](#) with a proof that if G is locally finite and quasi-transitive, then G_r is accessible. We deduce from this in [Section 7](#) that $N'(G_r)$ is exhaustive (cf. [Reduction 3](#)). In [Section 8](#) we infer that the separations in $N'(G_r)$ have finitely bounded order (cf. [Reduction 2](#)). This implies that G_r is tangle-accessible (cf. [Reduction 1](#)), which will allow us to prove [Theorem 5.5](#).

6. LOCAL COVERS ARE ACCESSIBLE

Our aim in this section is to establish [Theorem 6.2](#) below. Recall the definition of accessibility in graphs:

Definition 6.1 ([40]). A graph G is *accessible* if there exists an integer K such that every two ends of G are distinguished by a separation of order at most K .

Theorem 6.2. *Let G be a connected, locally finite, quasi-transitive graph, and let $r > 0$ be an integer. Then the r -local cover G_r of G is accessible.*

We already proved [Theorem 6.2](#) for the special cases of Cayley graphs ([Corollary 4.12](#)), and of arbitrary finite graphs by way of graph-decompositions ([Corollary 4.9](#)). These are the two cases that we would need for direct proofs of [Theorems 1](#) and [2](#), proceeding exactly as we shall below. However, [Theorem 6.2](#) is interesting in its own right, is considerably more general, and enables us to also prove [Theorem 5.5](#).

For a comprehensive proof of [Theorem 6.2](#), let us observe first that G_r is connected, locally finite, and quasi-transitive as soon as G is:

Lemma 6.3. *Let G be a connected, locally finite and quasi-transitive graph, and let $r > 0$ be an integer. Then the r -local cover G_r of G is connected, locally finite, and quasi-transitive.*

Proof. The r -local cover G_r is connected and locally finite by definition. To prove that G_r is quasi-transitive, let us show that lifts \hat{u} of u and \hat{v} of v lie in the same orbit of $\text{Aut}(G_r)$ as soon as $v = \varphi(u)$ for some $\varphi \in \text{Aut}(G)$. By [Lemma 3.11](#), and since local coverings are canonical, φ lifts to an automorphism $\hat{\varphi}$ of G_r . Then \hat{v} differs from $\hat{\varphi}(\hat{u})$ only by a deck transformation, so \hat{u} and \hat{v} lie in the same orbit of $\text{Aut}(G_r)$. \square

Our key ingredient for the proof of [Theorem 6.2](#) is the following result of Hamann:

Lemma 6.4 ([22, Corollary 3.3]). *Every locally finite quasi-transitive graph whose binary cycle space is generated by cycles of bounded length is accessible.*

Proof of [Theorem 6.2](#). By [Lemma 4.6](#), the binary cycle space of G_r is generated by its cycles of length at most r . The theorem now follows from [Lemma 6.4](#), which we may apply to G_r by [Lemma 6.3](#). \square

7. TOOLS AND LEMMAS FOR THE PROOF OF [THEOREM 1](#)

Throughout this section, G will be a connected locally finite infinite graph. Our main goal is to prove the following key lemma. Note that if G satisfies its premise, it must be accessible.

Lemma 7.1. *Let N be a nested set of tight finite-order separations of G , and suppose that N is Γ -canonical for some group Γ that acts quasi-transitively on G . Suppose further that there exists an integer K such that some Γ -canonical set $N_{\text{end}} \subseteq N$ of separations of order $< K$ distinguishes all the ends of G . Then N is exhaustive and induces a Γ -canonical tree-decomposition of G .*

7.1. Tools: Tree-decompositions and nested sets of separations. Let $\mathcal{T} = (T, \mathcal{V})$ be a tree-decomposition of G , and let Γ be a group acting on G . Recall from [Section 3.2](#) that every edge $t_1 t_2$ of T induces a separation $\{A_{t_1}, A_{t_2}\}$ of G , and that each of its two orientations (t_1, t_2) induces the oriented separation $\alpha_{\mathcal{T}}(t_1, t_2) = (A_{t_1}, A_{t_2})$ of G . If we orient distinct edges $t_1 t_2$ and $t_2 t_3$ of T each from t_i to t_{i+1} , then $\alpha_{\mathcal{T}}(t_1, t_2) \leq \alpha_{\mathcal{T}}(t_2, t_3)$ in the partial ordering \leq of oriented separations (see [Section 5.1](#)).

By definition of the sets A_{t_1} and A_{t_2} , the separations of G that \mathcal{T} induces in this way are pairwise nested. If \mathcal{T} is regular, they are proper separations. If \mathcal{T} is Γ -canonical for some group Γ as defined in [Section 3.3](#), then the nested set N of separations of G which \mathcal{T} induces is Γ -canonical as defined in [Section 5.1](#).

Conversely, let us try to construct from a given nested set N of proper separations of G a regular tree-decomposition of G that induces N . We begin by constructing, from N alone, the nodes of a potential decomposition tree. A *splitting star* of N is any set $\sigma \subseteq \vec{N}$ that consists of all the maximal separations of some consistent orientation O of N such that for every $\vec{r} \in O$ there exists $\vec{s} \in \sigma$ with $\vec{r} \leq \vec{s}$. Every separation in \vec{N} lies in at most one splitting star of N . If N is induced by a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G , as is our aim, then every $\vec{s} \in \vec{N}$ does lie in a splitting star: it will have the form $\vec{s} = \alpha_{\mathcal{T}}(t', t)$ for some oriented edge (t', t) of T , and \vec{s} will lie in splitting star of N formed by the separations $\alpha_{\mathcal{T}}(t'', t)$ where t'' varies over the neighbours of t in T .³

Taking our cue from the above observation, let us define a graph T from N as follows. Take the splitting stars of N as the nodes of T , and make two nodes adjacent if they contain, as splitting stars of N , distinct orientations of some element of N . For each splitting star $t \in V(T)$ let $V_t := \bigcap_{(A,B) \in t} B$, and put $\mathcal{V} = (V_t)_{t \in T}$. Let us say that N *induces* this pair (T, \mathcal{V}) .

As indicated, (T, \mathcal{V}) need not be a tree-decomposition of G when N is infinite. Indeed, while T is always acyclic it need not be connected, and the family \mathcal{V} might violate [\(T1\)](#) [[19](#), Example 4.9]. We shall mention some sufficient conditions for (T, \mathcal{V}) to be a tree-decomposition later.

If $\mathcal{T} := (T, \mathcal{V})$ is a tree-decomposition, it will induce the set N of separations we started with, in that $N = \{\alpha_{\mathcal{T}}(e) : e \in E(T)\}$ (as was our aim), with $\alpha_{\mathcal{T}}$ an order isomorphism between the oriented edges of T and \vec{N} [[19](#), Proof of Lemma 2.7]. If N is Γ -canonical for some group Γ , then \mathcal{T} will be a Γ -canonical tree-decomposition. Indeed, if N is Γ -canonical, then every $g \in \Gamma$ maps splitting stars of N to splitting stars of N . Since $g \cdot (A, B) = (g \cdot A, g \cdot B)$ for $g \in \Gamma$ and $(A, B) \in \vec{N}$, this action of Γ on T agrees with that in our definition of Γ -canonicity of (T, \mathcal{V}) in [Section 3.3](#).

A strictly increasing infinite sequence $(A_i, B_i)_{i \in \mathbb{N}}$ of oriented separations of G is *exhaustive* if $\bigcap_{i \in \mathbb{N}} B_i = \emptyset$. A set S of separations is *exhaustive* if every such sequence in \vec{S} is exhaustive. Any nested set of separations induced by a point-finite tree-decomposition of G is clearly exhaustive. We shall need the following converse:

Lemma 7.2 ([\[19](#), Lemma 2.7]). *Let N be a nested set of proper separations of G . If N is exhaustive, then the pair (T, \mathcal{V}) induced by N is a regular tree-decomposition of G . If a group Γ acts on G so that N is Γ -canonical, then this tree-decomposition (T, \mathcal{V}) is Γ -canonical.*⁴

The task of constructing tree-decompositions thus reduces to that of constructing exhaustive nested sets of separations. Let us collect a few lemmas that will help with this task. When $(A_i, B_i)_{i \in \mathbb{N}}$ and $(C_j, D_j)_{j \in \mathbb{N}}$ are sequences of separations such that for every i there exists j such that $(A_i, B_i) \leq (C_j, D_j)$, we say that $(C_j, D_j)_{j \in \mathbb{N}}$ *dominates* $(A_i, B_i)_{i \in \mathbb{N}}$.

³In general, when N is infinite, the elements of \vec{N} need not lie in a splitting star of N ; see [[12](#), [21](#)].

⁴The last sentence, which we already verified, is not explicitly stated in [[19](#)].

Lemma 7.3. *If $(A_i, B_i)_{i \in \mathbb{N}}$ and $(C_j, D_j)_{j \in \mathbb{N}}$ are sequences of separations such that $(C_j, D_j)_{j \in \mathbb{N}}$ dominates $(A_i, B_i)_{i \in \mathbb{N}}$, then $\bigcap_{i \in \mathbb{N}} B_i \supseteq \bigcap_{j \in \mathbb{N}} D_j$. In particular, if $(A_i, B_i)_{i \in \mathbb{N}}$ is exhaustive then so is $(C_j, D_j)_{j \in \mathbb{N}}$. \square*

All the separations from which we shall construct tree-decompositions will be tight, and hence proper.

Lemma 7.5 below, which is implicit in [19], shows that a nested set of tight separations is exhaustive if these have finitely bounded order. For its proof we need the following well-known fact about tight separations:

Lemma 7.4 ([40, Proposition 4.2]). *Every vertex of G lies in only finitely many separators of tight separations of order less than K , for each $K \in \mathbb{N}$.*

Lemma 7.5 ([19]). *Let K be an integer, and let $(A_i, B_i)_{i \in \mathbb{N}}$ be a strictly increasing sequence of tight separations of G , all of order $< K$. Then $(A_i, B_i)_{i \in \mathbb{N}}$ is exhaustive.*

Proof. By **Lemma 7.4**, every vertex of G lies in $A_i \cap B_i$ for only finitely many $i \in \mathbb{N}$. These sets $A_i \cap B_i$ are finite, so applying this to every vertex $v \in A_i \cap B_i$ shows that $(A_i \cap B_i) \cap (A_j \cap B_j) = \emptyset$ for all large enough j . We can thus find an infinite subsequence of $(A_i, B_i)_{i \in \mathbb{N}}$ whose separators are disjoint. But this implies that $\bigcap_i B_i = \emptyset$: otherwise G contains an $A_0 - \bigcap_i B_i$ path that meets all these infinitely many disjoint separators, a contradiction. Hence $(\vec{s}_i)_{i \in \mathbb{N}}$ is exhaustive. \square

7.2. Proof of Lemma 7.1 for 1-ended graphs. Let us now prove **Lemma 7.1** for 1-ended graphs. In fact, we show the following stronger result:

Lemma 7.6. *Assume that G has only one end. Let N be a nested set of tight separations of G , all of finite order. Suppose that N is Γ -canonical for some group Γ acting quasi-transitively on G . Then \vec{N} contains no strictly increasing infinite sequence.*

We shall need the following fact, proved in [14], about 1-ended, connected, locally finite graphs:

Lemma 7.7. *If G has only one end, then every finite-order separation of G has a finite and an infinite side.*

Our next lemma shows that, in a strictly increasing infinite sequence of separations, every separation is oriented towards its infinite side:

Lemma 7.8. *In any strictly increasing sequence $(A_i, B_i)_{i \in \mathbb{N}}$ of separations of G , all B_i are infinite.*

Proof. Suppose some B_i is finite; we may assume that $i = 0$. Since $B_0 \supseteq B_i$ for all i , all the B_i are finite. Since $(A_i, B_i)_{i \in \mathbb{N}}$ is a strictly increasing sequence, we have for each i either $A_i \subsetneq A_{i+1}$ or $B_i \supsetneq B_{i+1}$. As B_0 is finite, we have $B_i \supsetneq B_{i+1}$ for only finitely many i , so $A_i \subsetneq A_{i+1}$ for infinitely many i . For all these i the non-empty sets $A_{i+1} \setminus A_i$ are disjoint and contained in the finite set B_0 , a contradiction. \square

Recall that the set N in **Lemma 7.6**, which we are seeking to analyse, is Γ -canonical. The following lemma shows how separations with a finite side interact with their images under certain automorphisms:

Lemma 7.9. *Let $\{A, B\}$ be a separation of G with A finite. Consider any automorphism φ of G and a vertex $v \in A \cap B$ with a neighbour $w \in B \setminus A$ such that $\varphi(v) \in A \setminus B$. Then $\{A, B\}$ and $\{\varphi(A), \varphi(B)\}$ cross.*

Proof. Let us show first that neither $A \subseteq \varphi(A)$ nor $\varphi(A) \subseteq A$. To see this, it is enough to show $\varphi(A) \neq A$, since A is finite. If $\varphi(A) = A$, then $\varphi(v) \in A \setminus B$ is a neighbour of $\varphi(w) \in \varphi(B \setminus A) = B \setminus A$, which contradicts the fact that $\{A, B\}$ is a separation.

As B is infinite, we also have $\varphi(B) \not\subseteq A$. So if $\{A, B\}$ and $\{\varphi(A), \varphi(B)\}$ were nested, we would have $(A, B) \leq (\varphi(B), \varphi(A))$, and thus $B \supseteq \varphi(A)$. But $\varphi(v) \in \varphi(A) \setminus B$ by assumption, a contradiction. \square

We remark that [Lemma 7.9](#) has no analogue for finite graphs, or for separations with two infinite sides [\[14\]](#).

With these three lemmas at hand, we are now ready to prove [Lemma 7.6](#).

Proof of Lemma 7.6. Suppose, for a contradiction, that $(A_i, B_i)_{i \in \mathbb{N}}$ is a strictly increasing sequence in \vec{N} . Then $A_i \setminus B_i \subseteq A_j \setminus B_j$ for all $i < j$, and by [Lemmas 7.7](#) and [7.8](#) all the A_i are finite.

We first show that there exists $n \in \mathbb{N}$ such that $A_i \cap B_i \subseteq A_{i+1} \cap B_{i+1}$ for all $i \geq n$. If not, there are infinitely many i such that $A_i \cap B_i$ contains a vertex $v_i \in A_{i+1} \setminus B_{i+1}$. By our initial observation, these v_i are distinct for different i . Since Γ acts quasi-transitively on G , there exist $i < j$ among these with $v_i = g \cdot v_j$ for some $g \in \Gamma$. Then $\{A_j, B_j\}$, together with $v_j \in A_j \cap B_j$ and the automorphism φ of G defined by g , satisfies the premise of [Lemma 7.9](#), since $g \cdot v_j = v_i \in A_j \setminus B_j$ and $\{A_j, B_j\}$ is tight. Since N is Γ -canonical, applying [Lemma 7.9](#) yields a contradiction to N being nested.

We next show that there exists $m \geq n$ such that $A_i \cap B_i = A_{i+1} \cap B_{i+1}$ for all $i \geq m$. If not, then $A_i \cap B_i \subsetneq A_{i+1} \cap B_{i+1}$ for infinitely many $i \geq n$; pick a subsequence $(A_{i_j}, B_{i_j})_{j \in \mathbb{N}}$ with $i_0 = n$ witnessing this. Since the separations (A_{i_j}, B_{i_j}) are tight, every $G[A_{i_j} \setminus B_{i_j}]$ has a component K_{i_j} with neighbourhood $A_{i_j} \cap B_{i_j}$. Then $K_{i_j} \cap K_{i_{j+1}} = \emptyset$ for all j , since K_{i_j} is also a component of $G[A_{i_{j+1}} \setminus B_{i_{j+1}}]$ but distinct from $K_{i_{j+1}}$, since the two have different neighbourhoods. As G is connected, there exists $v \in A_n \cap B_n$. This v lies in $A_{i_j} \cap B_{i_j}$ for every j , and hence has a neighbour in K_{i_j} . As these neighbours are distinct, this contradicts our assumption that $v \in G$ has finite degree.

The (A_i, B_i) with $i \geq m$ form a strictly increasing sequence of separations with the same separator X . This is impossible since, as G is connected and locally finite, the graph $G - X$ has only finitely many components. \square

7.3. Proof of Lemma 7.1. In the remainder of this section we prove [Lemma 7.1](#). Let N , N_{end} , Γ and $K > 0$ be as stated in the lemma. Our goal is to show that N is exhaustive: that every strictly increasing sequence $(A_i, B_i)_{i \in \mathbb{N}}$ of separations in \vec{N} satisfies $\bigcap_i B_i = \emptyset$.

We already know from [Lemma 7.5](#) that N_{end} is exhaustive, since by assumption the separations in N_{end} have bounded order. If we can show that our sequence $(A_i, B_i)_{i \in \mathbb{N}}$ dominates a sequence in \vec{N}_{end} , it will also be exhaustive ([Lemma 7.3](#)). Essentially, this amounts to showing that our sequence is interlaced with a sequence in \vec{N}_{end} : that at most finitely many (A_i, B_i) have the same \leq -relationships to all the elements of \vec{N}_{end} . Our plan is to show that if there were infinitely many such (A_i, B_i) then these would, essentially, dissect the same 1-ended subgraph of G , which would contradict [Lemma 7.6](#).

In graph terms this means that, essentially, for the tree-decomposition $\mathcal{T}' = (T, \mathcal{V}')$ of G that N_{end} induces (see below), where $\mathcal{V}' = (V'_t)_{t \in T}$ say, there is no $t \in T$ such that infinitely many (A_i, B_i) differ only in how they cut up V'_t , because this would contradict [Lemma 7.6](#). This will not work directly, since the graphs $G[V'_t]$ need not satisfy the conditions on G laid down in the premise of [Lemma 7.6](#). However we shall be able to extend the V'_t to sets V_t that do satisfy these conditions, and then apply [Lemma 7.6](#) to the graphs $G[V_t]$ to complete our proof.

To get started, recall from [Lemmas 7.2](#) and [7.5](#) that N_{end} does induce the required tree-decomposition:

Lemma 7.10. *N_{end} is exhaustive, and the tree-decomposition \mathcal{T}' of G it induces is Γ -canonical.* \square

As N_{end} distinguishes all the ends of G , each part of \mathcal{T}' meets rays from at most one end of G infinitely.

Since N_{end} is Γ -canonical, the group Γ acts on the decomposition tree T of \mathcal{T}' (see Section 7.1). Since G is quasi-transitive under the action of Γ , there are only finitely many Γ -orbits of tight separations of G of order less than K [40, Proposition 4.2]. The action of Γ on T , therefore, is quasi-transitive too:

Lemma 7.11. Γ acts on $E(T)$ with finitely many orbits. \square

Let st be any edge of T . Since G is locally finite and $V'_{st} := V'_s \cap V'_t$ is finite, the graph $G - V'_{st}$ has only finitely many components. Hence there exists an integer $d_{st} \geq 1$ such that the d_{st} -ball $B_G(V'_{st}, d_{st})$ around V'_{st} (see Section 4.3) meets every component C of $G - V'_{st}$ in a set that includes all the shortest paths in C between neighbours of V'_{st} in C . (There are only finitely many such paths, since G is locally finite.) By Lemma 7.11, Γ has finitely many orbits on $E(T)$. Let us assume that the integers d_{st} were all chosen as small as possible; then they coincide for any two edges st in the same orbit. Thus, the maximum d over all d_{st} for $st \in E(T)$ is well-defined; let $V_t := V(B_G(V'_t, d))$ and $\mathcal{V} := (V_t)_{t \in T}$.

Lemma 7.12. $\mathcal{T} = (T, \mathcal{V})$ is a Γ -canonical tree-decomposition of G of finitely bounded adhesion such that the following statements hold for every $t \in T$:

- (i) $G[V_t]$ is connected;
- (ii) the stabiliser Γ_t of t under the action of Γ on T acts quasi-transitively on $G[V_t]$;
- (iii) $G[V_t]$ is either finite or 1-ended.

The proof of Lemma 7.12 closely follows [23]; details are given in [14].

Recall that our plan is to show that our sequence $(A_i, B_i)_{i \in \mathbb{N}}$ in \vec{N} dominates a suitable sequence in \vec{N}_{end} . To achieve this, it will suffice to show that no part V'_t of \mathcal{T}' contains all the separators from a final segment $(A_i, B_i)_{i \geq i_0}$ of our sequence. We will prove this by applying Lemma 7.6 to the corresponding part $G[V_t]$ of \mathcal{T} .

To apply Lemma 7.6 in this way, we need to show that a separation of G whose separator is contained in V'_t induces a tight separation of $G[V_t]$. It does, because the graph $G[V_t]$ behaves like a torso of $G[V'_t]$ (see [10, Ch.12.3] if desired):

Lemma 7.13. Let t be a node of T and $X \subseteq V'_t$. Then the map $C \mapsto C \cap G[V_t]$ is a bijection between the components of $G - X$ and those of $G[V_t] - X$. Moreover, $N_{G[V_t]}(C \cap G[V_t]) = N_G(C)$ for all such C .

Proof. Let C be any component of $G - X$. Since G is connected, C contains a neighbour of X . All the neighbours of X in C lie in V_t , since $N_G(X) \subseteq B_G(V'_t, 1) \subseteq B_G(V'_t, d) = V_t$. Hence C meets some component D of $G[V_t] - X$. We shall prove that D is the only component of $G[V_t] - X$ that C meets; then clearly $D = C \cap G[V_t]$, which establishes the desired bijection.

Since $G[V_t]$ is connected (Lemma 7.12), every component of $G[V_t] - X$ contains a neighbour of X . It thus suffices to show that all the neighbours u, v of X in C lie in the same component of $G[V_t] - X$ (our D). We prove this by showing that every shortest u - v path P in $G - X$ lies in $G[V_t]$. Note that $P \subseteq C$.

Let x be a neighbour of u in X , and y a neighbour of v in X , possibly $x = y$. Then $xuPvy$ is a concatenation of walks in G that either lie in $G[V'_t]$ or have their first and last vertex in V'_t and all their inner vertices (at least one) in the same component of $G - V'_t$. Let $Q = x' \dots y'$ be any subwalk of $xuPvy$ of the latter type if one exists (which will be a path or the entire walk $xuPvy$, in which case that is a cycle), and let C' be the component of $G - V'_t$ that contains $Q - x' - y'$. It will suffice for our proof of $P \subseteq G[V_t]$ to show that $Q \subseteq G[V_t]$.

As $X \subseteq V'_t$ and C' meets C (in $Q - x' - y' \neq \emptyset$, which lies in $P \subseteq C$), we have $C' \subseteq C$. Since \mathcal{T}' is a tree-decomposition, $V(C')$ is contained in $\bigcup_{t' \in T'} V'_{t'}$ for some component T' of $T - t$. Let $s \in T'$ be

the neighbour of t in that component. Then all the neighbours of C' in V'_t lie in $V'_{st} \subseteq V'_t$, and so the component of $G - V'_{st}$ that contains C' (which exists because $V'_{st} \subseteq V'_t$) in fact equals C' .

Thus, $Q - x' - y'$ is a path in a component of $G - V'_{st}$ between neighbours of V'_{st} (possibly identical). Since it is a subpath of P , it is also a shortest such path. Hence $Q \subseteq G[V_t]$ by the definition of V_t . \square

To complete our proof of [Lemma 7.1](#) we need one more definition. Let \vec{r}, \vec{s} be orientations of proper separations of G , such as those in \vec{N} . If $r = s$ we say that \vec{r} *points towards* \vec{s} if $\vec{r} = \vec{s}$. If $r \neq s$ we say that \vec{r} *points towards* \vec{s} (and s) if either $\vec{r} \leq \vec{s}$ or $\vec{r} \leq \vec{\bar{s}}$. Note that if $r \neq s$ are nested, they have unique orientations \vec{r} and $\vec{\bar{s}}$ that point towards each other, with $\vec{r} \leq \vec{s}$ and $\vec{\bar{r}} \geq \vec{\bar{s}}$. Also, since r and s are proper, we cannot have both $\vec{r} \leq \vec{s}$ and $\vec{r} \leq \vec{\bar{s}}$.

Completion of proof of Lemma 7.1. Recall that we have a sequence $(A_i, B_i)_{i \in \mathbb{N}}$ of separations in \vec{N} , which we want to show to be exhaustive. Let us abbreviate (A_i, B_i) as \vec{s}_i . The tools at our disposal are that $N_{\text{end}} \subseteq N$ is exhaustive ([Lemma 7.10](#)), our tree-decomposition of G into finite or 1-ended parts provided by [Lemma 7.12](#), and [Lemma 7.6](#) by which 1-ended graphs contain no infinite sequences such as $(\vec{s}_i)_{i \in \mathbb{N}}$ at all.

We need three more steps. We first establish a partition of \vec{N} defined by the splitting stars σ of N_{end} . Recall that these correspond to the nodes t of T (see [Section 7.1](#)); the corresponding partition class P_σ of \vec{N} will consist, very roughly, of those separations in \vec{N} that lie in σ or separate V_t . Next we show that for our proof that $(\vec{s}_i)_{i \in \mathbb{N}}$ is exhaustive it suffices to show that it contains separations from infinitely many of these partition classes P_σ . Finally, we check that $(\vec{s}_i)_{i \in \mathbb{N}}$ does meet infinitely many P_σ .

Since N_{end} is exhaustive, its splitting stars form a partition of \vec{N}_{end} : every element of \vec{N}_{end} lies in exactly one splitting star [[21](#), Lemma 2.4]. Let us extend this partition of \vec{N}_{end} to one of \vec{N} , as follows. For every splitting star σ of N_{end} let

$$P_\sigma := \{ \vec{s} \in \vec{N} : \text{all the elements of } \sigma \text{ point towards } \vec{s} \}.$$

Note that $P_\sigma \cap \vec{N}_{\text{end}} = \sigma$.

Let us show that these P_σ form a partition of \vec{N} . As the separations in N are nested, every $\vec{s} \in \vec{N}$ defines a consistent orientation of N_{end} :

$$O_{\vec{s}} := \{ \vec{r} \in \vec{N}_{\text{end}} : \vec{r} \text{ points towards } \vec{s} \};$$

let $\sigma_{\vec{s}}$ denote the set of its maximal elements. Since N_{end} is exhaustive and s is proper, $O_{\vec{s}}$ contains no infinite strictly increasing sequence. Hence every $\vec{r} \in O_{\vec{s}}$ lies below some maximal element of $O_{\vec{s}}$; thus, $\sigma_{\vec{s}}$ is a splitting star of N_{end} . So $P_{\sigma_{\vec{s}}}$ is defined, and it clearly contains \vec{s} . But \vec{s} cannot lie in any other P_σ . Indeed, consider a splitting star $\sigma \neq \sigma_{\vec{s}}$ of N_{end} . The orientation O of N_{end} of which σ is the set of maximal elements then orients some separation $r \in N_{\text{end}}$ differently from the way $O_{\vec{s}}$ does. Hence if $\vec{s} \in P_\sigma \cap P_{\sigma_{\vec{s}}}$ then both orientations of r lie below an element of σ or $\sigma_{\vec{s}}$, and hence both point towards \vec{s} . This is impossible, since s is proper.

Next we show that $(\vec{s}_i)_{i \in \mathbb{N}}$ is exhaustive if it meets infinitely many P_σ . By [Lemma 7.3](#) we may assume that the \vec{s}_i lie in distinct P_σ . For each $i \in \mathbb{N}$ let σ_i be such that $\vec{s}_i \in P_{\sigma_i}$. Then $O_i := O_{\vec{s}_i}$ is the consistent orientation of N_{end} of which σ_i is the set of maximal elements. Let $r_i \in N_{\text{end}}$ be a separation oriented differently by O_i and O_{i+1} . Then one of these, \vec{r}_i say, points towards \vec{s}_{i+1} , while the other, $\vec{\bar{r}}_i$, points towards \vec{s}_i . Since $\vec{s}_i < \vec{s}_{i+1}$, this is possible only if $\vec{s}_i < \vec{r}_i \leq \vec{s}_{i+1}$. (The first inequality is strict, because otherwise $\vec{\bar{r}}_i$ would point towards its own inverse $\vec{r}_i = \vec{\bar{s}}_i$, which is impossible.) Hence the sequence $(\vec{r}_i)_{i \in \mathbb{N}}$

is strictly increasing too, and is dominated by $(\vec{s}_i)_{i \in \mathbb{N}}$. Since $(\vec{r}_i)_{i \in \mathbb{N}}$ consists of separations from \vec{N}_{end} and is thus exhaustive by [Lemma 7.10](#), our sequence $(\vec{s}_i)_{i \in \mathbb{N}}$ is exhaustive by [Lemma 7.3](#).

It remains to check that $(\vec{s}_i)_{i \in \mathbb{N}}$ does meet infinitely many P_σ . To prove this we show that each P_σ contains only finitely many \vec{s}_i . Recall that $\mathcal{T}' = (T, \mathcal{V}')$ was the tree-decomposition of G induced by N_{end} , and that our group Γ of automorphisms of G acts on T . As shown in [Section 7.1](#), the splitting stars σ of N_{end} correspond to the nodes t of T . Given such a pair σ and t , let $N_t \subseteq N$ be the set of all separations with an orientation in P_σ . Note that the stabiliser Γ_t of t under the action of Γ on T consists precisely of those elements of Γ that fix σ as a set.

Let us show that N_t is Γ_t -canonical: that for every $g \in \Gamma_t$ and every $s \in N_t$ we have $g \cdot s \in N_t$. By definition of N_t and P_σ the latter means that every $\vec{r} \in \sigma$ points towards $g \cdot s$: that $s' := g \cdot s$ has an orientation $\vec{s}' \geq \vec{r}$. Let $\vec{r}' := g^{-1} \cdot \vec{r}$. As $g \in \Gamma_t$ we have $\vec{r}' \in \sigma$; thus, \vec{r}' also points towards s , say $\vec{r}' \leq \vec{s}$. But then $\vec{r}' = g \cdot \vec{r} \leq g \cdot \vec{s} =: \vec{s}'$ as required.

Let us turn N_t , a set of separations of G , into a Γ_t -canonical set M_t of tight separations of $G[V_t]$. By the definition of \mathcal{T}' , the part V'_t is equal to $\bigcap_{(C,D) \in \sigma} D$. As for every $\{A, B\} \in N_t$ and every $(C, D) \in \sigma$ either A or B is a subset of D , we thus have $A \cap B \subseteq V'_t$. Every oriented separation $(A, B) \in \vec{N}_t$ induces an oriented separation $(A_t, B_t) := (A \cap V_t, B \cap V_t)$ of $G[V_t]$. By [Lemma 7.13](#) one easily shows that these are distinct for different (A, B) , that they are tight because the (A, B) are tight, and that for $(A, B) < (C, D)$ in \vec{N}_t we have $(A_t, B_t) < (C_t, D_t)$. Let

$$M_t := \{\{A_t, B_t\} : \{A, B\} \in N_t\}.$$

M_t is a Γ_t -canonical set of tight separations of $G[V_t]$. If $G[V_t]$ is finite then so is M_t , and in particular \vec{M}_t contains no strictly increasing sequence. Otherwise $G[V_t]$ is a connected locally finite 1-ended graph that is quasi-transitive under the action of Γ_t , by [Lemma 7.12](#). Hence, by [Lemma 7.6](#), \vec{M}_t contains no strictly increasing infinite sequence. Since any strictly increasing sequence in \vec{N}_t induces a strictly increasing sequence in \vec{M}_t , by [Lemma 7.13](#) as noted earlier, there is thus no strictly increasing infinite sequence in \vec{N}_t either. This completes our proof that $P_\sigma \subseteq \vec{N}_t$ contains only finitely many \vec{s}_i , and hence that our sequence $(\vec{s}_i)_{i \in \mathbb{N}}$ is exhaustive. By [Lemma 7.2](#), N induces a Γ -canonical tree-decomposition of G , as required in the statement of [Lemma 7.1](#). \square

8. PROOF OF THE MAIN RESULT

In this section we combine all our lemmas to prove [Theorem 5.5](#) – the more general version of [Theorem 1](#) that also includes [Theorem 2](#) – and derive some corollaries of the proof. The unique graph-decomposition of the given graph G whose existence is asserted in [Theorem 5.5](#), the decomposition that displays its r -global structure (if it exists), was specified in [Definition 5.4](#). It is defined (via the projection $p_r : G_r \rightarrow G$) by the tree-decomposition \mathcal{T} of G_r that is induced by its canonical tree of tangles $N(G_r)$ from [Definition 5.2](#) if this exists.

We shall derive the existence of $N(G_r)$ from [Lemma 5.3](#), which requires G_r to be tangle-accessible. [Lemma 8.1](#) helps ensure this when G_r is accessible, which we already know from [Theorem 6.2](#).

Let us say that a set S of separations of a graph G is *point-finite* if for every vertex $v \in V(G)$ the set $S_v := \{\{A, B\} \in S : v \in A \cap B\}$ is finite. It is easy to see that a regular tree-decomposition is point-finite if (and only if) its induced nested set of separations is point-finite. Indeed, if a vertex of G lies in infinitely many parts then, by [\(T2\)](#) and [[10](#), Prop. 8.2.1], the decomposition tree has an infinite star or a ray such that

v lies in every part V_t corresponding to a node t in that star or ray. Then v also lies in all the separators that correspond to edges of that star or ray.

Lemma 8.1. *Let G be a connected locally finite graph, and let N be a nested set of tight finite-order separations of G . Let Γ be a group acting quasi-transitively on G so that N is Γ -canonical. Suppose that there exists a number $K \in \mathbb{N}$ such that some Γ -canonical set $N_{\text{end}} \subseteq N$ of separations of order $< K$ distinguishes all the ends of G .*

Then N is exhaustive and point-finite, and the order of the separations in N is finitely bounded. The tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G which N induces is regular, point-finite, Γ -canonical, and of finitely bounded adhesion. If G is 1-ended, then T is rayless.

For a proof of [Lemma 8.1](#) we need two observations.

Lemma 8.2. *Let G be a connected locally finite graph, and let N be a nested set of tight separations of G . If N is exhaustive, then N is point-finite.*

Proof. Let $v \in V(G)$, and suppose the set N_v of separations in N whose separator contains v is infinite. Since N is exhaustive, so is N_v . By [Lemma 7.2](#), N_v induces a regular tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G . If T contains a ray, then the forward-oriented edges of this ray define an infinite strictly increasing sequence of separations in N_v via $\alpha_{\mathcal{T}}$. All these separations have v in their separator, which contradicts the fact that N_v is exhaustive.

So T is rayless, and hence has a vertex t of infinite degree [[10](#), Prop.8.2.1]. Applying $\alpha_{\mathcal{T}}$ to the edges incident with t and oriented towards t yields an infinite set $\{(A_i, B_i) : i \in \mathbb{N}\} \subseteq \vec{N}_v$ such that $(A_i, B_i) < (B_j, A_j)$ for all $i \neq j$. All the sets $A_i \setminus B_i$ are disjoint, and v has a neighbour in every $A_i \setminus B_i$ since $v \in A_i \cap B_i$ and $\{A_i, B_i\}$ is tight. This contradicts the local finiteness of G . \square

In quasi-transitive graphs, canonical and point-finite sets of separations are again quasi-transitive under the action of the same group:

Lemma 8.3. *Let G be a connected graph, not necessarily locally finite, and let S be a point-finite set of finite-order separations of G . Let Γ be a group acting quasi-transitively on G so that S is Γ -canonical. Then Γ acts on S with finitely many orbits. In particular, the order of the separations in S is finitely bounded.*

Proof. As G is quasi-transitive under the action of Γ , there exists a finite set V' of vertices of G such that $\Gamma \cdot V' = V(G)$. Then $S \setminus \{\emptyset, V(G)\} = \Gamma \cdot S'$ where $S' = \bigcup_{v' \in V'} S_{v'}$. Since S is point-finite, all these $S_{v'}$ are finite, and hence so is S' . Thus, Γ acts on S with finitely many orbits, while $\{\emptyset, V(G)\}$, if it lies in S , forms an additional orbit in S . Since S' is finite, the order of the separations in S is finitely bounded. \square

Proof of [Lemma 8.1](#). By [Lemma 7.1](#), N is exhaustive. By [Lemma 8.2](#) it is point-finite, and by [Lemma 8.3](#) its separations have finitely bounded order. By [Lemma 7.2](#), N induces a regular Γ -canonical tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of G . Its adhesion sets are the separators of the separations in N , so it has finitely bounded adhesion. As remarked just before [Lemma 8.1](#), the point-finiteness of N implies that \mathcal{T} is point-finite, too.

If G is 1-ended then, by [Lemma 7.6](#), \vec{N} contains no infinite strictly increasing sequence. Any ray in T would define such a sequence via $\alpha_{\mathcal{T}}$, so T is rayless. \square

[Lemma 8.1](#) essentially strengthens accessibility to tangle-accessibility in our context. We need this for our proof of [Theorems 1 and 2](#), as outlined at the start of this section. More precisely, what we shall need

is [Theorem 3](#), which we will thus prove first. Our input N for [Lemma 8.1](#) in the proof of [Theorem 3](#) will come from the following non-trivial result:

Lemma 8.4 ([\[19, Theorem 6.6\]](#)). *Every connected, locally finite graph has a canonical tree of tangles consisting of relevant separations.*⁵

We can now prove [Theorem 3](#), which we restate:

Theorem 3. *Accessible locally finite, quasi-transitive graphs are even tangle-accessible.*

Proof. Let G be a graph as stated. Since it is quasi-transitive, it has only finitely many isomorphism types of components, so we may assume G to be connected. Let $N := N'(G)$ be the tree of tangles given by [Lemma 8.4](#).

By [Lemma 5.1](#), N is a set of tight finite-order separations of G . Since ends define tangles, and these are distinguished by exactly the separations that distinguish those ends ([Section 5.1](#)), N distinguishes the ends of G efficiently.

Since G is accessible, the order of the separations of G needed to distinguish its ends is finitely bounded. As N is canonical and distinguishes the ends efficiently, it thus has a canonical subset N_{end} of separations of finitely bounded order that achieves this too – for example, its subset of all efficient end-distinguishers. By [Lemma 8.1](#), then, the separations in N have finitely bounded order too. Since N distinguishes the set of all tangles of G , it thus witnesses that G is tangle-accessible. \square

Note that the tree of tangles $N'(G)$ given by [Lemma 8.4](#), which we just used in the proof of [Theorem 3](#), exists for arbitrary connected locally finite graphs, not just quasi-transitive ones. In general it may have limit points, in which case it will not define a tree-decomposition. But even when G is tangle-accessible and $N'(G)$ does define a tree-decomposition, it need not coincide with our canonical tree of tangles $N(G)$ from [Definition 5.2](#).

Turning now to our canonical tree of tangles $N(G)$ from [Definition 5.2](#), recall that in order to ensure its existence we needed that G is tangle-accessible ([Lemma 5.3](#)). By [Theorem 3](#), this now follows already from classical accessibility, which we shall obtain from [Theorem 6.2](#).

Our next lemma records that, in this case, $N(G)$ induces the desired tree-decomposition:

Lemma 8.5. *Let G be a connected, locally finite, quasi-transitive, accessible graph. Then $N(G)$ is defined and induces a canonical and point-finite regular tree-decomposition of G of finitely bounded adhesion. If G is 1-ended, the decomposition tree is rayless.*

Proof. By [Theorem 3](#) and [Lemma 5.3](#), $N(G)$ exists, is canonical, and consists of relevant (and therefore tight) separations of finitely bounded order. Since $N(G)$ distinguishes all the ends of G , it induces the desired tree-decomposition by [Lemma 8.1](#). \square

Summing everything up, we can now prove [Theorem 5.5](#):

Proof of Theorem 5.5. Let G and r be as stated in the theorem. The r -local cover G_r of G is connected, locally finite, and quasi-transitive by [Lemma 6.3](#), and accessible by [Theorem 6.2](#). By [Lemma 8.5](#), $N(G_r)$ is defined and induces a canonical and point-finite regular tree-decomposition \mathcal{T} of G_r . As $p_r : G_r \rightarrow G$ is also canonical (see [Section 4.2](#)), \mathcal{T} defines a canonical point-finite graph-decomposition of G ([Theorem 3.13](#)). By [Definition 5.4](#), this is the unique graph-decomposition of G that displays its r -global structure. \square

⁵The relevance of the separations is not explicitly stated in [\[19, Theorem 6.6\]](#) but immediate from its proof.

Proof of Theorem 1. **Theorem 1** is the special case of **Theorem 5.5** where G is finite. □

Proof of Theorem 2. Since $\text{Cay}(\Gamma, S)$ is locally finite and transitive, **Theorem 5.5** implies **Theorem 2**. □

REFERENCES

- [1] A. Amit and N. Linial, *Random graph coverings I: General theory and connectivity*, *Combinatorica* **22** (2002), 1–18.
- [2] N. Bowler, *personal communication*, 2021.
- [3] J. Carmesin, *Local k -separators*. Manuscript 2021.
- [4] ———, *Local 2-separators*, *Journal of Combinatorial Theory, Series B* **156** (2022; arXiv:2008.03032), 101–144.
- [5] J. Carmesin, R. Diestel, M. Hamann, and F. Hundertmark, *k -Blocks: a connectivity invariant for graphs*, *SIAM Journal on Discrete Mathematics* **28** (2014; arXiv:1305.4557), 1876–1891.
- [6] ———, *Canonical tree-decompositions of finite graphs I. Existence and algorithms*, *Journal of Combinatorial Theory, Series B* **116** (2016; arXiv:1406.3797), 1–24.
- [7] J. Carmesin, R. Diestel, F. Hundertmark, and M. Stein, *Connectivity and tree structure in finite graphs*, *Combinatorica* **34** (2014; arXiv:1105.1611), no. 1, 1–35.
- [8] J. Carmesin and J.P. Gollin, *Canonical tree-decompositions of a graph that display its k -blocks*, *Journal of Combinatorial Theory, Series B* **122** (2017; arXiv:1506.02904), 1–20.
- [9] J. Carmesin and J. Kurkofka, *Entanglements*, (2022; arXiv:2205.11488). Submitted.
- [10] R. Diestel, *Graph theory* (5th edition), Springer-Verlag, 2017.
Electronic edition available at <http://diestel-graph-theory.com/>.
- [11] ———, *Ends and tangles*, *Abhandlungen Math. Sem. Univ. Hamburg* **87**, Special issue in memory of Rudolf Halin (2017; arXiv:1510.04050), 223–244.
- [12] ———, *Abstract separation systems*, *Order* **35** (2018; arXiv:1406.3797), 157–170.
- [13] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, *Combinatorica* **39** (2019; arXiv:1110.6207), no. 1, 37–75.
- [14] R. Diestel, R.W. Jacobs, P. Knappe, and J. Kurkofka, *Canonical graph decompositions via coverings, extended version*, 2022. arXiv:2207.04855v3.
- [15] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, *Journal of Combinatorial Theory, Series B* **87** (2003), 197–206.
- [16] ———, *Graph minor hierarchies*, *Discrete Applied Mathematics* **145** (2005), 167–182.
- [17] D.Ž. Djoković, *Automorphisms of graphs and coverings*, *Journal of Combinatorial Theory, Series B* **16** (1974), no. 3, 243–247.
- [18] M.J. Dunwoody, *The accessibility of finitely presented groups*, *Inventiones Mathematicae* **81** (1985), no. 3, 449–457.
- [19] C. Elbracht, J. Kneip, and M. Teegen, *Trees of tangles in infinite separation systems*, *Math. Proc. Camb. Phil. Soc.* (2021; arXiv:2005.12122), 1–31.
- [20] A. Georgakopoulos, *On covers of graphs by Cayley graphs*, *European Journal of Combinatorics* **64** (2017; arXiv:1504.00173), 57–65.
- [21] J.P. Gollin and J. Kneip, *Representations of infinite tree sets*, *Order* (2020; arXiv:1908.10327).
- [22] M. Hamann, *Accessibility in transitive graphs*, *Combinatorica* **38** (2018; arXiv:1404.7677), no. 4, 847–859.
- [23] M. Hamann, F. Lehner, B. Miraftab, and T. Rühmann, *A Stallings-type theorem for quasi-transitive graphs*, *Journal of Combinatorial Theory, Series B* **157** (2022; arXiv:1812.06312), 40–69.
- [24] A. Hatcher, *Algebraic Topology*, Cambridge Univ. Press, 2002.
- [25] R.W. Jacobs and P. Knappe, *Local coverings of plane graphs*. In preparation.
- [26] ———, *Locally chordal graphs*. In preparation.
- [27] R.W. Jacobs, P. Knappe, and J. Kurkofka, *A hierarchy of local connectivity*. In preparation.
- [28] ———, *Local blocks and how to distinguish them*. In preparation.
- [29] ———, *Local blueprints of graphs*. In preparation.
- [30] K. Jänich, *Topology*, Springer Berlin Heidelberg, 1997.
- [31] C. Löh, *Geometric Group Theory*, Springer, 2017.
- [32] W. Magnus, A. Karrass, and D. Solitar, *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*, Courier Corporation, 2004.

- [33] J. Milnor, *A note on curvature and fundamental group*, Journal of Differential geometry **2** (1968), no. 1, 1–7.
- [34] N. Robertson and P.D. Seymour, *Graph minors I–XX*, Journal of Combinatorial Theory, Series B (1983).
- [35] ———, *Graph minors. X. Obstructions to tree-decomposition*, Journal of Combinatorial Theory, Series B **52** (1991), 153–190.
- [36] A. Scott, *Graphs of large chromatic number*, ICM 2022.
- [37] J.R. Stallings, *Group theory and three-dimensional manifolds*, Yale University Press, 1972.
- [38] ———, *Topology of finite graphs*, Inventiones mathematicae **71** (1983), no. 3, 551–565.
- [39] A.S. Švarc, *A volume invariant of coverings*, Dokl. akad. nauk sssr (ns), 1955, pp. 32–34.
- [40] C. Thomassen and W. Woess, *Vertex-transitive graphs and accessibility*, Journal of Combinatorial Theory, Series B **58** (1993), 248–268.
- [41] C.T.C. Wall, *Pairs of relative cohomological dimension one*, Journal of Pure and Applied Algebra **1** (1971), no. 2, 141–154.