

A Stallings' type theorem for quasi-transitive graphs

Matthias Hamann*

Alfréd Rényi Institute of Mathematics
Budapest, Hungary

Florian Lehner †

Mathematics Institute, University of Warwick
Coventry, UK

Babak Miraftab

Department of Mathematics, University of Hamburg
Hamburg, Germany

Tim Rühmann

Department of Mathematics, University of Hamburg
Hamburg, Germany

April 15, 2019

Abstract

We consider infinite connected quasi-transitive locally finite graphs and show that every such graph with more than one end is a tree amalgamation of two other such graphs. This can be seen as a graph-theoretical version of Stallings' splitting theorem for multi-ended finitely generated groups and indeed it implies this theorem. It will also lead to a characterisation of accessible graphs. We obtain applications of our results for hyperbolic graphs, planar graphs and graphs without any thick end. The application for planar graphs answers a question of Mohar in the affirmative.

*Supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement n° 617747 and through the Heisenberg-Programme of the Deutsche Forschungsgemeinschaft (DFG Grant HA 8257/1-1).

†Supported by the Austrian Science Fund (FWF), grant J 3850-N32

1 Introduction

Stallings [18] proved that finitely generated groups with more than one end are either a free product with amalgamation over a finite subgroup or an HNN-extension over a finite subgroup. The main aim of this paper is to obtain an analogue of Stallings' theorem for quasi-transitive graphs. The obvious obstacle for this is that free products with amalgamations and HNN-extensions are group theoretical concepts. So in order to obtain a graph-theoretical analogue, we first need to find a graph-theoretical analogue of free products with amalgamations and HNN-extensions. The proposed notation by Mohar [14] are tree amalgamations and indeed we will prove the following theorem. (We refer to Section 5 for the definition of tree amalgamations.)

Theorem 1.1. *Every connected quasi-transitive locally finite graph with more than one end is a non-trivial tree amalgamation of finite adhesion of two connected quasi-transitive locally finite graphs.*

On the other side, we can ask if we start with finite or one-ended connected quasi-transitive locally finite graphs and do iterated tree amalgamations of finite adhesion, what class of graphs do we end up with? In the case of finitely generated groups, the answer is the class of accessible groups (by definition). Thomassen and Woess [20] defined edge-accessibility for graphs¹: a quasi-transitive locally finite graph is *edge-accessible* if there is some $n \in \mathbb{N}$ such that every two ends can be separated by at most n edges. They showed in [20] that a finitely generated group is accessible if and only if each of its locally finite Cayley graphs is edge-accessible. We will show that tree amalgamations and edge-accessibility fit well together in that we prove that the above described class of graphs we obtain is the class of edge-accessible connected quasi-transitive locally finite graphs.

In 1988, Mohar [14] asked whether tree amalgamations are powerful enough to yield a classification of infinitely-ended transitive planar graphs in terms of finite and one-ended infinite planar transitive graphs. Our theorems enable us to answer his question in the affirmative for quasi-transitive graphs because Dunwoody [9] proved that they are edge-accessible, see Section 7.3.

Additionally, we obtain as a corollary Stallings' theorem, see Section 7.1, and a new characterisation of quasi-transitive locally finite graphs that are quasi-isometric to trees, see Section 7.2. In Section 7.4 we apply our theorems to hyperbolic graphs and show that a quasi-transitive locally finite graph is hyperbolic if and only if it is obtained by iterated tree amalgamations starting with finite or one-ended hyperbolic quasi-transitive locally finite graphs.

Our main tool to prove Theorem 1.1 are canonical tree-decompositions. While some proofs of Stallings' theorem are build on edge separators and their structure trees, see e. g. Dunwoody [6], it turns out that tree-decompositions and vertex separators fit well together with tree amalgamations. Due to the similar natures of structure trees and tree-decompositions, it is not surprising that some results that we prove here (Propositions 4.7 and 4.8) have also been proved for structure trees, see e. g. Thomassen and Woess [20] and Möller [15, 16].

¹They call it accessible instead of edge-accessible, but we will reserve the notion of accessibility to a direct translation of accessibility of groups, see Section 6.

2 Preliminaries

We follow the general notations of [5] unless stated otherwise. In the following we will state the most important definitions for convenience.

Let $G = (V(G), E(G))$ be a graph. A *geodesic* is a shortest path between two vertices. A *ray* is a one-way infinite path, the infinite subpaths of a ray are its *tails*. Two rays are *equivalent* if there exists no finite vertex set separating them eventually, i. e. two rays are equivalent if they have tails contained in the same component of $G - S$ for every finite set S of vertices. The equivalence classes of rays in a graph are its *ends*. The *degree* of an end is the maximum number of disjoint rays in that end, if it exists. If that maximum does not exist, we say that this end has *infinite degree* and call it *thick*. An end with finite degree is called *thin*. An end ω is *captured* by a set X of vertices if every ray of ω has infinite intersection with X and it *lives* in X if every ray of ω has a tail in X .

Let $X \subseteq V(G)$. Let G' be the graph with vertex set $(V(G) \setminus X) \cup \{v_X\}$, where v_X is a new vertex, and edges between $u, v \in V(G) \setminus X$ if and only if $uv \in E(G)$ and v_X is adjacent to precisely those vertices $y \in V(G) \setminus X$ that have a neighbour in X . We call G' the *contraction* of X in G and we say that we *contracted* X . Since edges are just vertex sets of size 2, the definition carries over to edges.

Let Γ be a group acting on G and let $X \subseteq V(G)$. The (*setwise*) *stabilizer* of X with respect to Γ is the set

$$\Gamma_X := \{g \in \Gamma \mid g(x) \in X \text{ for all } x \in X\}.$$

An *orbit* of Γ (or a Γ -*orbit*) is a set $\{g(x) \mid g \in \Gamma\}$ for some $x \in V(G)$. We say Γ acts *transitively* on G if $V(G)$ is one Γ -orbit and Γ acts *quasi-transitively* on G if $V(G)$ consists of finitely many Γ -orbits.

3 Canonical tree-decompositions

In this section we will look at our main tool for our proofs: tree-decompositions. A *tree-decomposition* of a graph G is a pair (T, \mathcal{V}) where T is a tree and $\mathcal{V} = (V_t)_{t \in V(T)}$ is a family of vertex sets of G such that the following holds:

(T1) $V(G) = \bigcup_{t \in V(T)} V_t$.

(T2) For every edge $e \in E(G)$ there is a $t \in V(T)$ such that V_t contains both vertices that are incident with e .

(T3) $V_{t_1} \cap V_{t_2} \subseteq V_{t_3}$ whenever t_3 lies on the $t_1 - t_2$ path in T .

The sets V_t are called the *parts* of (T, \mathcal{V}) and the vertices of the *decomposition tree* T are its *nodes*. The sets $V_{t_1} \cap V_{t_2}$ with $t_1 t_2 \in E(T)$ are the *adhesion sets* of the tree-decomposition. We say that (T, \mathcal{V}) has *finite adhesion* if all adhesion sets are finite.

Remark 3.1. Let $t_1 t_2$ be an edge of the decomposition tree T of a tree-decomposition (T, \mathcal{V}) . For $i = 1, 2$, let T_i be the component of $T - t_1 t_2$ that contains t_i . It follows from (T3) that $V_{t_1} \cap V_{t_2}$ separates the vertices in $\bigcup_{t \in T_1} V_t$ from those in $\bigcup_{t \in T_2} V_t$.

We say (T, \mathcal{V}) *distinguishes* two ends ω_1 and ω_2 if there is a finite adhesion set $V_{t_1} \cap V_{t_2}$ such that one end lives in $\bigcup_{t \in T_1} V_t$ and the other lives in $\bigcup_{t \in T_2} V_t$, where T_i is the maximal subtree of $T - t_1 t_2$ containing t_i . It distinguishes them *efficiently* if no vertex set in G of smaller size than $V_{t_1} \cap V_{t_2}$ separates them. For $k \in \mathbb{N}$, two ends of G are *k-distinguishable* if there is a set of k vertices of G that separates them.

Let Γ be a group acting on G . If every $\gamma \in \Gamma$ maps parts of (T, \mathcal{V}) to parts and thereby induces an automorphism of T we say that (T, \mathcal{V}) is Γ -*invariant*.

The following theorem by Carmesin et al. will be the main result we are building on.

Theorem 3.2. [2] *Let G be a locally finite graph, let Γ be a group acting on G and let $k \in \mathbb{N}$. Then there is a Γ -invariant tree-decomposition of G of adhesion at most k that efficiently distinguishes all k -distinguishable ends. \square*

4 Basic tree-decompositions

The aim of this section is first to modify the tree-decomposition of Theorem 3.2 and then to prove some properties of the newly obtained tree-decomposition, in particular, where the tree-decomposition captures the ends of the graph. Our first step in modifying the tree-decomposition of Theorem 3.2 will be to make all adhesion sets connected while keeping the action of Γ on (T, \mathcal{V}) .

Proposition 4.1. *Let Γ be a group acting on a locally finite graph G and let $(T, \mathcal{V}) := (T, (V_t)_{t \in V(T)})$ be a Γ -invariant tree-decomposition of G of finite adhesion. Then there is a Γ -invariant tree-decomposition $(T, \mathcal{V}') := (T, (V'_t)_{t \in V(T)})$ of G such that every adhesion set of (T, \mathcal{V}') is finite and connected and such that $V_t \subseteq V'_t$ for every $t \in V(T)$.*

Proof. Let u and v be two vertices of an adhesion set of (T, \mathcal{V}) . Assume that \mathcal{P}_{uv} is the set of all geodesics between u and v and assume that V_{uv} is the set of all vertices of G that lie on the paths of \mathcal{P}_{uv} . For a part V_t let V'_t be the union of V_t with all sets V_{uv} where u and v lie in an adhesion set contained in V_t . Let $\mathcal{V}' := \{V'_t \mid t \in V(T)\}$. We claim that (T, \mathcal{V}') is a tree-decomposition. By construction it has the desired properties, i. e. every adhesion set is connected and $V_t \subseteq V'_t$ and, since G is locally finite and since the adhesion sets of (T, \mathcal{V}) are finite, every adhesion set of (T, \mathcal{V}') is finite. Since we made no choices when adding all possible geodesics to the adhesion sets, Γ still acts on (T, \mathcal{V}') .

As every element of \mathcal{V}' is a superset of some element of \mathcal{V} , we just have to verify (T3) to see that (T, \mathcal{V}') is a tree-decomposition. To see this, let $x \in V'_{t_1} \cap V'_{t_2}$ for $t_1, t_2 \in V(T)$ and let t_3 be on the t_1 - t_2 path s_1, \dots, s_n in T with $s_1 = t_1$ and $s_n = t_2$. If $x \in V_{t_1} \cap V_{t_2}$, then we have $x \in V_{t_3} \subseteq V'_{t_3}$ as (T, \mathcal{V}) is a tree-decomposition. If $x \in (V'_{t_1} \setminus V_{t_1}) \cap V_{t_2}$, then it lies on a geodesic P between two vertices x_1, x_2 of an adhesion set of (T, \mathcal{V}) in V_{t_1} . Since every adhesion set $V_{s_i} \cap V_{s_{i+1}}$ separates V_{s_1} from V_{s_n} and since $x \in V_{t_2}$, the path P must pass through $V_{s_i} \cap V_{s_{i+1}}$. Thus, either P contains two vertices u, v of $V_{s_i} \cap V_{s_{i+1}}$ such that x lies on the u - v subpath P' of P , or x lies in $V_{s_i} \cap V_{s_{i+1}}$. In the first case, we added P' to the adhesion set $V_{s_i} \cap V_{s_{i+1}}$ because P' is a geodesic with its end vertices in $V_{s_i} \cap V_{s_{i+1}}$. Thus, in both cases x lies in $V_{s_i} \cap V_{s_{i+1}}$ and thus in V'_{t_3} . If $x \in (V'_{t_1} \setminus V_{t_1}) \cap (V'_{t_2} \setminus V_{t_2})$, let $t_4 \in V(T)$ with $x \in V_{t_4}$. By the previous

case, x lies in V'_t for every t on the t_1 - t_4 or t_2 - t_4 paths in T . Since T is a tree, these cover the path s_1, \dots, s_n and hence $x \in V'_{t_3}$. This proves that (T, \mathcal{V}') is a tree-decomposition. \square

We call a tree-decomposition of a graph G *connected* if all parts induce connected subgraphs of G .

The step to make the adhesion sets connected is just an intermediate step for us as we aim for connected parts, i. e. we aim for connected tree-decompositions. The next lemma ensures that in connected graphs all parts are connected if all adhesion sets are connected.

Lemma 4.2. *If all adhesion sets of a tree-decomposition of a connected graph are connected, then the tree-decomposition is connected.*

Proof. Let G be a graph and let (T, \mathcal{V}) be a tree-decomposition of G all of whose adhesion sets are connected. Let u and w be two vertices of V_t for some $t \in V(T)$. Since G is connected, there is a path $P = p_1, \dots, p_n$ with $p_1 = u$ and $p_n = w$. We choose P with as few vertices outside of V_t as possible. Let us suppose that P leaves V_t . Let $p_i \in V_t$ such that $p_{i+1} \notin V_t$ and let p_j be the first vertex of P after p_i that lies in V_t . As $p_n = w \in V_t$ we know that such a vertex always exists. Let $t' \in V(T)$ such that $p_{i+1} \in V_{t'}$. Then the adhesion set $V_t \cap V_{t'}$, where s is the neighbour of t on the t - t' path in T , separates V_t from p_{i+1} . Hence, the definition of a tree-decomposition implies that p_j must lie in $V_t \cap V_s$, too. But then we can replace the subpath of P between p_i and p_j by a path in $V_t \cap V_s$. The resulting walk contains a path between u and w with fewer vertices outside of V_t than P . This contradiction shows that all vertices of P lie in V_t and hence $G[V_t]$ is connected. \square

Most of the time we do not need the full strength of Theorem 3.2 in that it suffices to consider Γ -invariant tree-decompositions with few Γ -orbits that still distinguish some ends.

Let Γ be a group acting on a connected locally finite graph G with at least two ends. A Γ -invariant tree-decomposition (T, \mathcal{V}) of G is a *basic tree-decomposition (with respect to Γ)* if it has the following properties:

- (i) (T, \mathcal{V}) distinguishes at least two ends.
- (ii) Every adhesion set of (T, \mathcal{V}) is finite.
- (iii) Γ acts on (T, \mathcal{V}) with precisely one orbit on $E(T)$.

If it is clear from the context which group we consider, we just say that (T, \mathcal{V}) is a basic tree-decomposition of G . It follows from Theorem 3.2 that basic tree-decompositions always exist:

Corollary 4.3. *Let Γ be a group acting on a locally finite graph G with at least two ends. Then there is a basic tree-decomposition (T, \mathcal{V}) for G .*

Proof. By Theorem 3.2, we find a Γ -invariant tree-decomposition (T, \mathcal{V}) of bounded adhesion that separates some ends. Let tt' be an edge of T such that $V_t \cap V_{t'}$ separates some ends. Let $E_{tt'}$ be the orbit of tt' , i. e. the set $\{g(tt') \mid g \in \Gamma\}$, and let T' be obtained from T by contracting each component C of $T - E_{tt'}$ to a single vertex t_C . We set $V_{t_C} := \bigcup_{s \in C} V_s$ and set \mathcal{V}' be the

set of those sets V_{t_C} . It is easy to see that (T', \mathcal{V}') is a basic tree-decomposition with respect to Γ : the only non-trivial requirement is that (T', \mathcal{V}') distinguishes at least two ends. But this follows from the fact that $V_t \cap V_{t'}$ separates two ends. \square

Let us combine our results on connected and basic tree-decompositions.

Corollary 4.4. *Let Γ be a group acting on a connected locally finite graph G with at least two ends. Then the following hold.*

- (i) *There is a connected basic tree-decomposition of G with respect to Γ .*
- (ii) *If $(T, (V_t)_{t \in V(T)})$ is a basic tree-decomposition of G with respect to Γ , then there is a connected basic tree-decomposition $(T, (V'_t)_{t \in V(T)})$ of G with respect to Γ such that $V_t \subseteq V'_t$ for every $t \in V(T)$.*

Proof. By Corollary 4.3, there is a basic tree-decomposition of G . Having a basic tree-decomposition $(T, (V_t)_{t \in V(T)})$, Proposition 4.1 and Lemma 4.2 imply the existence of a connected basic tree-decomposition $(T, (V'_t)_{t \in V(T)})$ with $V_t \subseteq V'_t$ for every $t \in V(T)$. \square

Now we investigate some of the connections between the graphs and the parts of any of the connected basic tree-decompositions. We start by showing that these tree-decompositions behave well with respect to the class of quasi-transitive graphs.

Proposition 4.5. *Let Γ be a group acting quasi-transitively on a connected locally finite graph G with at least two ends and let (T, \mathcal{V}) be a connected basic tree-decomposition of G . Then for each part $V_t \in \mathcal{V}$ its stabilizer Γ_{V_t} acts quasi-transitively on $G[V_t]$.*

Proof. If $u \in V_t$ does not lie in any adhesion set, then none of its images $v \in V_t$ under elements of Γ lie in an adhesion set. Hence, if $\gamma \in \Gamma$ maps u to v , it must fix V_t setwise, as it acts on (T, \mathcal{V}) , so it lies in the stabilizer of V_t . Thus, the intersection of V_t with the Γ -orbit of u is the Γ_{V_t} -orbit of u .

Now we consider the vertices in an adhesion set $V_t \cap V_{t'}$. Let $V_t \cap V_s$ be another adhesion set. As (T, \mathcal{V}) is basic, there exists $\gamma \in \Gamma$ that maps $V_t \cap V_{t'}$ to $V_t \cap V_s$. If γ stabilizes V_t , all vertices of $V_t \cap V_s$ lie in Γ_{V_t} -orbits of the vertices of $V_t \cap V_{t'}$. Let us assume that γ does not stabilize V_t and let $V_t \cap V_{s'}$ be another adhesion set such that the element $\gamma' \in \Gamma$ that maps $V_t \cap V_{t'}$ to $V_t \cap V_{s'}$ does not stabilize V_t . Then $\gamma'\gamma^{-1}$ maps $V_t \cap V_s$ to $V_t \cap V_{s'}$ and stabilizes V_t . We conclude that the number of Γ_{V_t} -orbits of vertices in adhesion sets of V_t is at most twice the number of Γ -orbits of vertices in adhesion sets of V_t . \square

Subtrees of connected basic tree-decompositions that contain a common adhesion set cannot be too large as the following lemma shows.

Lemma 4.6. *Let Γ be a group acting quasi-transitively on a connected locally finite graph G with at least two ends and let (T, \mathcal{V}) be a connected basic tree-decomposition of G with respect to Γ . For an adhesion set X let T_X be the maximal subtree of T such that $X \subseteq V_t$ for all $t \in V(T_X)$. Then the diameter of T_X is at most 2.*

Proof. Suppose the diameter of T_X is at least 3. We have $V_t \cap V_{t'} = X$ for every $tt' \in E(T_X)$ since X is contained in every adhesion set $V_t \cap V_{t'}$ and since they all have the same size. Let $R = \dots t_0 t_1 \dots$ be a maximal path in T_X . We shall show that R is a double ray.

Let us suppose that t_{i+3} is the last vertex on R . As (T, \mathcal{V}) is basic, we find $\gamma \in \Gamma$ such that $\gamma(t_i t_{i+1}) = t_{i+2} t_{i+3}$. Note that γ fixes $X = V_{t_i} \cap V_{t_{i+1}} = V_{t_{i+2}} \cap V_{t_{i+3}}$ setwise. If $\gamma(t_i) = t_{i+2}$, then $\gamma(t_{i+2})$ is a neighbour of t_{i+3} distinct from t_{i+2} that contains X , a contradiction to the choice of i . If $\gamma(t_i) = t_{i+3}$, then γ fixes the edge $t_{i+1} t_{i+2}$ but neither of its incident vertices. Let $\gamma' \in \Gamma$ map $t_{i+1} t_{i+2}$ to $t_{i+2} t_{i+3}$. Note that γ' fixes X setwise, too. Then either γ' or $\gamma'\gamma$ maps t_i to a neighbour of t_{i+3} distinct from t_{i+2} . This is again a contradiction, which shows that R has no last vertex. Analogously, R has no first vertex. So it is a double ray.

Note that the part of some node of T_X contains X properly as $G = X$ is finite otherwise. But as Γ acts transitively on $E(T)$, we have at most two Γ -orbits on $V(T)$. Hence infinitely many parts of R contain X properly. Thus and since each V_{t_i} is connected, one vertex of X must have infinitely many neighbours. This contradiction to local finiteness shows the assertion. \square

Our next result is a characterisation of the finite parts of a connected basic tree-decomposition.

Proposition 4.7. *Let Γ be a group acting quasi-transitively on a connected locally finite graph G with at least two ends and let (T, \mathcal{V}) be a connected basic tree-decomposition of G . Then the degree of a node $t \in V(T)$ is finite if and only if V_t is finite.*

Proof. Note that each vertex lies in only finitely many adhesion sets as we only have one orbit of adhesion sets and as G is locally finite. So if V_t is finite, then the degree of t is finite, too.

Now let us assume that the degree of t is finite. Let U be a subset of V_t that consists of one vertex from each Γ_{V_t} -orbit that meets V_t . By Proposition 4.5 the set U is finite. The vertices in U have bounded distance to the union W of all adhesion sets in V_t . As they meet all Γ_{V_t} -orbits and Γ_{V_t} fixes W setwise, all vertices in V_t have bounded distance to W . Note that W is finite as t has finite degree. Since G is locally finite, V_t must be finite. \square

Let (T, \mathcal{V}) be a tree-decomposition of a graph G . We say that an end η of T captures an end ω of G if for every ray $R = t_1, t_2, \dots$ in η the union $\bigcup_{i \in \mathbb{N}} V_{t_i}$ captures ω and a node of T captures ω if its part does so.

Let us now investigate where the ends of G lie in (T, \mathcal{V}) .

Proposition 4.8. *Let G be a graph and let (T, \mathcal{V}) be a connected tree-decomposition of G such that the maximum size of its adhesion sets is at most $k \in \mathbb{N}$. Then the following holds.*

- (i) *Each end of G is captured either by an end or by a node of T .*
- (ii) *Every thick end of G is captured by a node of T .*
- (iii) *Every end of T captures a unique thin end of G , which has degree at most k .*

(iv) Assume that Γ acts quasi-transitively on G and that (T, \mathcal{V}) is Γ -invariant with only finitely many Γ -orbits on $E(T)$. Every end of G that is captured by a node $t \in V(T)$ corresponds to a unique end of $G[V_t]$.²

Proof. Let ω be an end of G and let Q, R be two rays in ω . For an edge $st \in E(T)$ let T_s and T_t be the subtrees of $T - st$ with $s \in V(T_s)$ and $t \in V(T_t)$. If the ray Q has all but finitely many vertices in $\bigcup_{x \in V(T_s)} V_x$ and R has all but finitely many vertices in $\bigcup_{x \in V(T_t)} V_x$ or vice versa, then we have a contradiction as Q and R cannot lie in the same end if they have tails that are separated by the finite vertex set $V_s \cap V_t$. We now orient the edge st from s to t if Q and R lie in $\bigcup_{x \in V(T_t)} V_x$ eventually and we orient it from t to s if the rays lie in $\bigcup_{x \in V(T_s)} V_x$ eventually. Obviously, every node of T has at most one outgoing edge. Let t_Q, t_R be nodes of T such that the first vertex of Q lies in V_{t_Q} , and the first vertex of R lies in V_{t_R} , and let P_Q and P_R be the maximal (perhaps infinite) directed paths in our orientation of T that start at t_Q and t_R , respectively. Note that if P_Q and P_R meet at a vertex, they continue in the same way. Thus, if they meet, they either end at a common vertex or have a common infinite subpath. We shall show that P_Q and P_R meet. Let P be the t_Q - t_R path in T . Then there is a unique sink x on it as every node of T has at most one outgoing edge. This sink is a common node of P_Q and P_R . If P_Q and P_R end at a node, this node captures ω and if they share a common infinite subpath, this is a ray whose end captures ω . We proved (i).

Now let us assume that ω has degree at least $k + 1$. Then there are $k + 1$ pairwise disjoint rays R_1, \dots, R_{k+1} in ω . Let t_i, P_i be a node and a path of T defined for R_i as we defined t_R and P_R for the ray R . By an easy induction, we can extend the above argument that P_Q and P_R meet to obtain that all P_1, \dots, P_{k+1} have a common node x . Let us suppose that ω is captured by an end η of T . Let y be the node of T that is adjacent to x and that separates x and η . Then all rays R_i must contain a vertex of $V_x \cap V_y$. This is not possible as $V_x \cap V_y$ contains at most k vertices and the rays R_i are disjoint. This contradiction shows (ii) and the second part of (iii).

Let R, Q be two rays that lie in ends of G that are captured by the same end η of T . With the notations P_Q, P_R as above, the intersection $P_Q \cap P_R$ is a ray in ω . As G is locally finite and (T, \mathcal{V}) is a connected tree-decomposition, there are infinitely many disjoint paths between Q and R and thus, they are equivalent and lie in the same end of G . This proves (iii).

To prove (iv), let us assume that Γ acts quasi-transitively on G and has finitely many orbits on the edges of the decomposition tree T . Let ω be an end of G that is captured by a node $t \in V(T)$ and let R be a ray in ω that starts at a vertex in V_t . Since V_t captures ω , there are infinitely many vertices of V_t on R . Whenever R leaves V_t through an adhesion set, it must reenter it through the same adhesion set by Remark 3.1. We replace every such subpath P , where the end vertices of P lie in a common adhesion set and the inner vertices of P lie outside of V_t , by a geodesic in $G[V_t]$ between the end vertices of P . We end up with a walk W with the same starting vertex as R . We shall see that W contains a one-way infinite path. First, we recursively delete closed subwalks of W to end up with a path R' . Since G is locally finite and R meets V_t infinitely often, R contains vertices of V_t that are arbitrarily far away from the starting vertex

²This shall mean that for every end ω of G that is captured by $t \in V(T)$ there is a unique end ω_t of $G[V_t]$ with $\omega_t \subseteq \omega$.

of R . As we only took geodesics to replace the subpaths of R that were outside of V_t and as Γ acts on (T, \mathcal{V}) with only finitely many orbits on the edges of T , these replacement paths have a bounded length. Hence, W eventually leaves every ball of finite diameter around its starting vertex. This implies that R' is a ray. Obviously, R and R' are equivalent. Thus, $G[V_t]$ contains a ray in ω . Let ω_t be the end of $G[V_t]$ that contains R' and let Q be a ray in ω_t . Since no finite separator can separate Q and R' in $G[V_t]$, the rays are also equivalent in G . Thus, we have shown $\omega_t \subseteq \omega$.

Let ω'_t be an end in $G[V_t]$ different from ω_t , let S be a finite subset of V_t that separates ω_t from ω'_t , and let P be a path in G connecting vertices in different components of $G[V_t] - S$. As before, whenever P leaves V_t through an adhesion set, it must reenter it through the same adhesion set by Remark 3.1. We again replace every such subpath, where the end vertices lie in a common adhesion set and the inner vertices lie outside of V_t , by a geodesic in $G[V_t]$ to obtain a walk P' in $G[V_t]$. Since P and P' have the same endpoints and P' must meet S , we know that P either contains a vertex in S , or it contains a vertex in an adhesion set which meets S . Let S' be the set containing all vertices of S and all vertices contained in adhesion sets that meet S . There are only finitely many orbits of vertices in adhesion sets, hence there is an upper bound on the diameter of the adhesion sets. Since S is finite and G is locally finite, this implies that S' is finite. By definition, there is no path in $G - S'$ connecting vertices in different components of $G[V_t] - S$. In particular, S' separates every ray in ω_t from every ray in ω'_t , and hence (iv) holds. \square

5 Tree amalgamations

In this section, we prove our main result, Theorem 1.1. But before we move on to that proof, we first have to state some definitions, in particular, the main definition: tree amalgamations, a notion introduced by Mohar [14]. After we stated those definitions, we compare tree amalgamations and connected basic tree-decompositions.

For the definition of tree amalgamations, let G_1 and G_2 be graphs. Let $(S_k^i)_{k \in I_i}$ be a family of subsets of $V(G_i)$. Assume that all sets S_k^i have the same cardinality and that the index sets I_1 and I_2 are disjoint. For all $k \in I_1$ and $\ell \in I_2$, let $\phi_{k\ell}: S_k^1 \rightarrow S_\ell^2$ be a bijection and let $\phi_{\ell k} = \phi_{k\ell}^{-1}$. We call the maps $\phi_{k\ell}$ and $\phi_{\ell k}$ *bonding maps*.

Let T be a $(|I_1|, |I_2|)$ -*semiregular* tree, that is, a tree in which for the canonical bipartition $\{V_1, V_2\}$ of $V(T)$ the vertices in V_i all have degree $|I_i|$. Denote by $D(T)$ the set obtained from the edge set of T by replacing every edge xy by two directed edges \vec{xy} and \vec{yx} . For a directed edge $\vec{e} = \vec{xy} \in D(T)$, we denote by $\overleftarrow{e} = \vec{yx}$ the edge with the reversed orientation. Let $f: D(T) \rightarrow I_1 \cup I_2$ be a labelling, such that for every $t \in V_i$, the labels of edges starting at t are in bijection to I_i .

For every $i \in \{1, 2\}$ and for every $t \in V_i$, take a copy G_t of the graph G_i . Denote by S_k^t the corresponding copies of S_k^i in $V(G_t)$. Let us take the disjoint union of the graphs G_t for all $t \in V(T)$. For every edge $\vec{e} = \vec{st}$ with $f(\vec{e}) = k$ and $f(\overleftarrow{e}) = \ell$ we identify each vertex x in the copy of S_k^s with the vertex $\phi_{k\ell}(x)$ in S_ℓ^t . Note that this does not depend on the orientation we pick for \vec{e} , since $\phi_{\ell k} = \phi_{k\ell}^{-1}$. The resulting graph is called the *tree amalgamation* of the graphs G_1

and G_2 over the *connecting tree* T and is denoted by $G_1 * G_2$ or by $G_1 *_T G_2$ if we want to specify the tree.

In the context of tree amalgamations the sets S_k^i are called the *adhesion sets* of the tree amalgamation. More specifically, the sets S_k^1 are the adhesion sets of G_1 and the sets S_k^2 are the adhesion sets of G_2 . If the adhesion sets of a tree amalgamation are finite, then this tree amalgamation has *finite adhesion*. We call a tree amalgamation $G_1 *_T G_2$ *trivial* if for some $t \in V(T)$ the canonical map that maps the vertices $x \in V(G_t)$ to the vertices of $G_1 *_T G_2$ that is obtained from x by all the identifications is a bijection. Note that if the tree amalgamation has finite adhesion, it is trivial if $V(G_i)$ is the only adhesion set of G_i and $|I_i| = 1$ for some $i \in \{1, 2\}$.

We remark that the map described in the definition of a trivial tree amalgamation does not induce a graph isomorphism $G_t \rightarrow G_1 *_T G_2$: it is a bijection $V(G_t) \rightarrow V(G_1 *_T G_2)$ but need not induce a bijection $E(G_t) \rightarrow E(G_1 *_T G_2)$.

The *identification length* of a vertex $x \in V(G_1 *_T G_2)$ is the diameter of the subtree T' of T induced by all nodes t for which a vertex of G_t is identified with x . The *identification length* of the tree amalgamation is the supremum of the identification lengths of its vertices. The tree amalgamation has *finite identification length* if the identification length is finite.

We remark that in Mohar's definition of a tree amalgamation [14] the identification length is always at most 2. But apart from this, our definition is equivalent to his.

It is worth noting that every tree amalgamation gives rise to a tree decomposition in the following sense.

Remark 5.1. Let G be a graph. If G is a tree amalgamation $G_1 *_T G_2$ of finite adhesion, then there is a naturally defined tree-decomposition of G : for $t \in V(T)$ let V_t be the set obtained from $V(G_t)$ after all identifications in $G_1 * G_2$. Set $\mathcal{V} := \{V_t \mid t \in V(T)\}$. Obviously, all vertices of G lie in $\bigcup_{t \in V(T)} V_t$ and for each edge there is some $V_t \in \mathcal{V}$ containing it. Property (T3) of a tree-decomposition is satisfied as the copies G_i^v are arranged in a treelike way and as identifications to obtain a vertex take place in subtrees of T . So (T, \mathcal{V}) is a tree-decomposition. If $G_1 *_T G_2$ has finite adhesion, so does (T, \mathcal{V}) . If the tree amalgamation is non-trivial, then T has at least two ends and so does G . Also, (T, \mathcal{V}) distinguishes two ends of G : those that are captured by ends of T .

So far, the tree amalgamations do not interact with any group actions on G_1 and G_2 . In particular, it is easy to construct a tree amalgamation of two quasi-transitive graphs that is not quasi-transitive: e. g. take as G_1 a double ray and as G_2 a finite non-trivial graph. Let G_1 have precisely two adhesion sets and G_2 at least two, all of size 1. The tree amalgamation $G_1 * G_2$ is not quasi-transitive.

In the following, we describe some conditions on tree amalgamations which will ensure that tree amalgamations of quasi-transitive graphs are again quasi-transitive, see Lemma 5.3.

Let Γ_i be a group acting on G_i for $i = 1, 2$, let $t \in V_i$, let $\gamma \in \Gamma_i$ and let $j \in \{1, 2\} \setminus \{i\}$. We say that the tree amalgamation *respects* γ , if there is a permutation π of I_i such that for every $k \in I_i$ there is $\ell \in I_j$ such that

$$\phi_{k\ell} = \phi_{\pi(k)\ell} \circ \gamma |_{S_k} .$$

Note that this in particular implies that $\gamma(S_k) = S_{\pi(k)}$. The tree amalgamation respects Γ_i if it respects every $\gamma \in \Gamma_i$.

Let $k \in I_i$ and let $\ell, \ell' \in I_j$. We call the bonding maps from k to ℓ and ℓ' *consistent* if there is $\gamma \in \Gamma_j$ such that

$$\phi_{k\ell} = \gamma \circ \phi_{k\ell'}.$$

We say that the bonding maps between two sets $J_1 \subseteq I_1$ and $J_2 \subseteq I_2$ are *consistent*, if they are consistent for any $i \in \{1, 2\}$, $k \in J_i$, and $\ell, \ell' \in J_j$.

We say that the tree amalgamation $G_1 * G_2$ is of *Type 1 respecting the actions of Γ_1 and Γ_2* or $(G_1, \Gamma_1) * (G_2, \Gamma_2)$ is a *tree amalgamation of Type 1* for short if the following holds:

- (i) The tree amalgamation respects Γ_1 and Γ_2 .
- (ii) The bonding maps between I_1 and I_2 are consistent.

We say that the tree amalgamation $G_1 * G_2$ is of *Type 2 respecting the actions of Γ_1 and Γ_2* or $(G_1, \Gamma_1) * (G_2, \Gamma_2)$ is a *tree amalgamation of Type 2* for short if the following holds:

- (o) $G_1 = G_2 =: G$, $\Gamma_1 = \Gamma_2 =: \Gamma$, and $I_1 = I_2 =: I$,³ and there is $J \subseteq I$ such that $f(\vec{e}) \in J$, if and only if $f(\bar{e}) \notin J$.
- (i) The tree amalgamation respects Γ .
- (ii) The bonding maps between J and $I \setminus J$ are consistent.

In this second case also say that $G_1 * G_2 = G * G$ is a *tree amalgamation of G with itself*.

We say that $G_1 * G_2$ is a tree amalgamation *respecting the actions of Γ_1 and Γ_2* if it is of either Type 1 or Type 2 respecting the actions Γ_1 and Γ_2 and we speak about the *tree amalgamation* $(G_1, \Gamma_1) * (G_2, \Gamma_2)$.

Note that conditions (i) and (ii) in both cases do not depend on the specific labelling of the tree. This is no coincidence. In fact we will show that any two legal labellings of $D(T)$ give isomorphic tree amalgamations, see Lemma 5.3. Furthermore, any $\gamma \in \Gamma_i$ (interpreted as an isomorphism between parts of two such tree amalgamations) can be extended to an isomorphism of the tree amalgamations, which also implies that the tree amalgamations obtained this way are always quasi-transitive.

Before we turn to the proof of these facts, we need some notation. A *legally labelled star centred at V_i* is a function ℓ from I_i to I_j . If the tree amalgamation is of Type 2, we further require that $\ell(k) \in J$ if and only if $k \notin J$. Informally, think of this as a star whose labels on directed edges could appear on a subtree of T induced by a vertex $t \in V_i$ and its neighbours: for \vec{e} with label k , the value $\ell(k)$ tells us the label of \bar{e} .

An isomorphism of two legally labelled stars ℓ, ℓ' is a triple $(\gamma, \pi, (\gamma_k)_{k \in I_i})$ consisting of some $\gamma \in \Gamma_i$, a permutation π of I_i , and a family $(\gamma_k)_{k \in I_i}$ of elements of Γ_j such that for every $k \in I_i$

$$\phi_{k, \ell(k)} = \gamma_k \circ \phi_{\pi(k) \ell'(\pi(k))} \circ \gamma|_{S_k}.$$

³Technically this is not allowed, in particular since for the definition of $\phi_{k\ell}$ we needed I_1 and I_2 to be disjoint. These technicalities can be easily dealt with by an appropriate notion of isomorphism the details of which we leave to the reader.

In our interpretation of legally labelled stars as subtrees of T , this corresponds to an isomorphism of the corresponding subgraphs of the tree amalgamation.

Proposition 5.2. *Let ℓ, ℓ' be two legally labelled stars with respect to a tree amalgamation $(G_1, \Gamma_1) *_{T'} (G_2, \Gamma_2)$ centred at V_i and let $\gamma \in \Gamma_i$. Then γ extends to an isomorphism $(\gamma', \pi, (\gamma_k)_{k \in I_i})$ of ℓ and ℓ' . Furthermore, if we are given $\tilde{k}, \tilde{k}' \in I_i$ and $\tilde{\gamma}_k \in \Gamma_j$ such that*

$$\phi_{\tilde{k}, \ell(\tilde{k})} = \tilde{\gamma}_k \circ \phi_{\tilde{k}', \ell'(\tilde{k}')} \circ \gamma|_{S_k},$$

then we can choose $\pi(\tilde{k}) = \tilde{k}'$ and $\gamma_{\tilde{k}} = \tilde{\gamma}_k$.

Proof. Since the tree amalgamation respects γ , there are π and $\bar{\ell}: I_i \rightarrow I_j$ such that

$$\phi_{k\bar{\ell}(k)} = \phi_{\pi(k)\bar{\ell}(k)} \circ \gamma|_{S_k}.$$

Let $\gamma_k \in \Gamma_j$ be such that $\phi_{k\bar{\ell}(k)} = \gamma_k \circ \phi_{k\bar{\ell}(k)}$, and let $\gamma'_k \in \Gamma_j$ be such that $\phi_{\pi(k)\bar{\ell}(k)} = \gamma'_k \circ \phi_{\pi(k)\bar{\ell}(k)}$. These exist by (ii); for Type 2 recall that by the definition of legally labelled stars $k \in J$ if and only if $\bar{\ell}(k) \notin J$. Now clearly

$$\phi_{k\bar{\ell}(k)} = \gamma_k \circ \gamma'_k \circ \phi_{\pi(k)\bar{\ell}(k)} \circ \gamma|_{S_k},$$

thus showing that the two stars are isomorphic.

For the second part, let $(\gamma', \pi, (\gamma_k)_{k \in I_i})$ be an isomorphism between ℓ and ℓ' . Let $\tilde{k}'' = \pi^{-1}(\tilde{k}')$. Define $\tau(\tilde{k}) = \tilde{k}'$ and $\tau(\tilde{k}'') = \pi(\tilde{k})$. Let $\delta_{\tilde{k}} = \tilde{\gamma}_k$ and let

$$\delta_{\tilde{k}''} = \gamma_{\tilde{k}''} \circ \tilde{\gamma}_k^{-1} \circ \gamma_{\tilde{k}}.$$

For the remaining $k \in I_i$, let $\tau(k) = \pi(k)$ and $\delta_k = \gamma_k$. It is straightforward to check that γ, τ , and $(\delta_k)_{k \in I_i}$ define an isomorphism between ℓ and ℓ' with the desired properties. \square

Lemma 5.3. *Let G_1 and G_2 be connected locally finite graphs and let Γ_i be a group acting quasi-transitively on G_i for $i = 1, 2$. Then the tree amalgamation $(G_1, \Gamma_1) *_{T'} (G_2, \Gamma_2)$ is quasi-transitive and independent (up to isomorphism) of the particular labelling of T .*

Proof. Let T and T' be two labelled trees giving rise to tree amalgamations $G = (G_1, \Gamma_1) *_{T'} (G_2, \Gamma_2)$ and $G = (G_1, \Gamma_1) *_{T'} (G_2, \Gamma_2)$, respectively, such that the adhesion sets as well as the bonding maps for both tree amalgamations are the same. Let $t \in V(T)$ and let $t' \in V(T')$ such that G_t and $G_{t'}$ are both isomorphic to G_i . Let $\gamma_t \in \Gamma_i$. We claim that there is an isomorphism $\bar{\gamma}: G \rightarrow G'$ such that

$$\bar{\gamma}|_{G_t} = \text{id}_{t'} \circ \gamma_t \circ \text{id}_t^{-1},$$

where id_t and $\text{id}_{t'}$ denote the canonical isomorphisms from G_i to G_t and $G_{t'}$ respectively. Clearly, the lemma follows from this claim.

For the proof of the claim define the star around $s \in V(T)$ by the map ℓ_s mapping k to the label of \vec{e}_k , where \vec{e}_k is the unique edge with label k starting at s . By Proposition 5.2, there are a bijection $\pi: N(t) \rightarrow N(t')$ and a family $(\gamma_s \in \Gamma_j)_{s \in N(t)}$ which extend γ_t to an isomorphism of the stars around t and t' . Iteratively apply Proposition 5.2 to vertices at distance $n = \{1, 2, 3, \dots\}$ from t . We obtain an isomorphism $\pi: T \rightarrow T'$ and maps $\gamma_s \in \Gamma_i$ for each $s \in V_i$ such

that the restriction of π to s and its neighbours and the corresponding maps γ_x form an isomorphism between the stars at s and $\pi(s)$.

For $v \in V(G_s)$, define $\bar{\gamma}(v) = \text{id}_{\pi(s)} \circ \gamma_s \circ \text{id}_s^{-1}(v)$. Note that for any edge ss' the two concurring definitions given for vertices of G_s and $G_{s'}$ that get identified for the tree amalgamation coincide. Hence $\bar{\gamma}$ is well defined, and since it obviously maps edges to edges and non-edges to non-edges, it is the desired isomorphism. \square

A closer inspection of the proof of Lemma 5.3 together with Remark 5.1 shows that tree amalgamations respecting the actions of quasi-transitive groups give rise to basic tree-decompositions of $(G_1, \Gamma_1) * (G_2, \Gamma_2)$. The following lemma shows that the converse also holds, that is, basic tree-decompositions of quasi-transitive graphs give rise to tree amalgamations respecting the actions of some quasi-transitive group on the parts.

Lemma 5.4. *Let Γ be a group acting quasi-transitively on a connected locally finite graph G and let (T, \mathcal{V}) be a connected basic tree-decomposition of G with respect to Γ . Then one of the following holds.*

- (1) *There are $V_t, V_{t'} \in \mathcal{V}$ such that G is a non-trivial tree amalgamation*

$$G[V_t] *_T G[V_{t'}]$$

of Type 1 respecting the actions of the stabilisers of $G[V_t]$ and $G[V_{t'}]$ in Γ .

- (2) *There is $V_t \in \mathcal{V}$ such that G is a non-trivial tree amalgamation*

$$G[V_t] *_T G[V_t]$$

of Type 2 respecting the actions of the stabiliser of $G[V_t]$ in Γ .

Proof. Choose an oriented edge $\vec{e}_0 \in D(T)$. We say that $\vec{e} \in D(T)$ is positively oriented, if there is $\gamma \in \Gamma$ mapping \vec{e}_0 to \vec{e} . Otherwise we say that \vec{e} is negatively oriented. If Γ contains an element that reverses an edge of T , then let Γ' be the subgroup preserving the bipartition of T . This subgroup has index 2, and still acts quasi-transitively on G and transitively on edges of T . Hence we can without loss of generality assume that no element of Γ swaps the endpoints of an edge, and thus every edge is either positively or negatively oriented, but not both.

Let s and t be the start and end point of \vec{e}_0 respectively. Let $(\vec{e}_k)_{k \in K}$ be the positively oriented edges starting at s and let $(\vec{e}_\ell)_{\ell \in L}$ be the negatively oriented edges starting at t . Without loss of generality, assume that K and L are disjoint, and that $\vec{e}_0 = \vec{e}_{k_0} = \vec{e}_{\ell_0}$. For every $k \in K$ pick $\gamma_k \in \Gamma$ which maps \vec{e}_0 to \vec{e}_k (with $\gamma_{k_0} = \text{id}$). For every $\ell \in L$ pick $\gamma_\ell \in \Gamma$ which maps \vec{e}_0 to \vec{e}_ℓ (with $\gamma_{\ell_0} = \text{id}$). If there is an element of Γ that maps s to t , then fix such an automorphism γ_{st} , and for $k \in K, \ell \in L$ let $\gamma'_k = \gamma_k \circ \gamma_{st}$ and $\gamma'_\ell = \gamma_\ell \circ \gamma_{st}^{-1}$.

Note that e_0 can be mapped to any edge incident to e_0 by a unique element of the form γ_k or γ'_k for some $k \in K \cup L$. For an arbitrary edge e , let e' be the first edge of the path connecting e to e_0 . If $\gamma_{e'} \in \Gamma$ maps e_0 to e' , then by the above remark there is a unique element δ_e of the form γ_k or γ'_k such that $\gamma_e \circ \delta_e$ maps e_0 to e . Use this to inductively construct (starting from $\delta_{e_0} = \text{id}$) for each $e \in E(T)$ an automorphism $\gamma_e \in \Gamma$ such that $\gamma_e(e_0) = e$. Let \vec{e} be the

orientation of e pointing away from e_0 if $e \neq e_0$ and $\vec{e} = \vec{e}_0$ otherwise. Define the label $f(\vec{e})$ to be the unique $k \in K \cup L$ such that the δ_e from above is γ_k or γ'_k . Note that $k \in K$ if and only if \vec{e} is positively oriented. In this case define $f(\vec{e}) = \ell_0$, otherwise define $f(\vec{e}) = k_0$.

The following observation will be useful later. Let v be a vertex of T , and let \vec{e} be the first edge of the path from v to e_0 (in case v is s or t this is an orientation of e_0). Let $\Delta_v = \{\delta_f \mid v \in f, f \neq e\}$.

- If all edges starting at v are positively (resp. negatively) oriented, then $\Delta_v = \{\gamma_k \mid k_0 \neq k \in K\}$ (resp. $\Delta_v = \{\gamma_\ell \mid \ell_0 \neq \ell \in L\}$).
- Otherwise, if \vec{e} is positively (resp. negatively) oriented, then $\Delta_v = \{\gamma_k, \gamma'_\ell \mid k_0 \neq k \in K, \ell \in \ell\}$ (resp. $\Delta_v = \{\gamma'_k, \gamma_\ell \mid k \in K, \ell_0 \neq \ell \in L\}$).

In particular, taking into account the label of \vec{e} , in the first case the edges starting at v are labelled bijectively by K (resp. L), while in the second case they are labelled bijectively by $K \cup L$.

Next we show how this labelling defines a tree amalgamation. First assume there is no automorphism $\gamma \in \Gamma$ mapping s to t . Then all positively oriented edges must point from V_1 to V_2 , where $V_1 \cup V_2$ is the bipartition of T with $s \in V_1$ —this corresponds to the first case in the above observation. Let G_1 be isomorphic to $G[V_s]$, and let G_2 be isomorphic to $G[V_t]$. Let id_s and id_t be the respective isomorphisms.

For the definition of the adhesion sets let $I_1 = K$ and $I_2 = L$. For $k \in K$ let t_k be the endpoint of \vec{e}_k and define $S_k = \text{id}_s^{-1}(V_s \cap V_{t_k})$. Similarly, for $\ell \in L$, let s_ℓ be the endpoint of \vec{e}_ℓ and define $S_\ell = \text{id}_t^{-1}(V_t \cap V_{s_\ell})$. Finally, define the adhesion maps by $\phi_{k\ell} = \text{id}_t^{-1} \circ \gamma_\ell \circ \gamma_k^{-1} \circ \text{id}_s \mid_{S_k}$.

The labels of directed edges starting at each vertex are in bijection to K or L depending on whether the vertex is in V_1 or V_2 . Hence the above information together with the labelling defines a tree amalgamation. If Γ_i is a group acting on G_i in the same way as the setwise stabiliser (in Γ) of G_s , of G_t acts on G_s , on G_t respectively, then it is straightforward to verify that this tree amalgamation is of Type 1 respecting the actions. Note that the possible replacement of Γ by Γ' changes neither Γ_1 nor Γ_2 .

It only remains to show that the tree amalgamation is isomorphic to G . Let e_v be the first edge on the path from $v \in V(T)$ to e_0 . If $v \in V_1$, then set $\text{id}_v = \gamma_{e_v} \circ \text{id}_s$. Otherwise set $\text{id}_v = \gamma_{e_v} \circ \text{id}_t$. It is easy to verify that for an edge $e = uv$ with labels $f(\vec{e}) = k, f(\vec{e}) = \ell$ we have that $\text{id}_v^{-1} \circ \text{id}_u = \phi_{k\ell}$, and this clearly shows that the tree amalgamation is isomorphic to G .

The proof in the case where there is γ_{st} mapping s to t is very similar to the first case. Define G_1 and G_2 as before, but make sure that $\gamma_{st} \circ \text{id}_s = \text{id}_t$. This ensures that the actions Γ_1 on G_1 and Γ_2 on G_2 are the same, hence we can without loss of generality assume that $G_1 = G_2$ and $\Gamma_1 = \Gamma_2$.

Set $I_1 = I_2 = K \cup L$ and let $J = K$. Recall that $f(\vec{e}) \in K$ if and only if $f(\vec{e}) \in L$. Since we can map s to t , there are positively and negatively oriented edges starting at each vertex, hence the labels of edges starting at any vertex are in bijection with $K \cup L$. Hence (o) for tree amalgamations of Type 2 holds. Define the adhesion sets and adhesion maps exactly as above (but note that all adhesion sets end up in the same graph since $G_1 = G_2$). This gives a tree amalgamation of Type 2 by construction which is isomorphic to G by the same argument as above. \square

Now we are ready to prove the main result of this section, the graph-theoretical analogue of Stallings' theorem, Theorem 1.1. We are proving a slightly stronger version than the one we stated in the introduction.

Theorem 5.5. *Let Γ be a group acting quasi-transitively on a locally finite graph G with more than one end. Then there are subgraphs G_1, G_2 of G and groups Γ_1, Γ_2 acting quasi-transitively on G_1, G_2 , respectively, such that G is a non-trivial tree amalgamation $(G_1, \Gamma_1) * (G_2, \Gamma_2)$ of finite adhesion and finite identification length.*

Furthermore, Γ_i can be chosen to be the setwise stabiliser of G_i in Γ .

Proof. By Corollary 4.4, G has a connected basic tree-decomposition (T, \mathcal{V}) . Using Lemma 5.4, G is a non-trivial tree amalgamation $(G_1, \Gamma_1) *_{T'} (G_2, \Gamma_2)$, where Γ_i is the setwise stabiliser of G_i in Γ for $i = 1, 2$. Proposition 4.5 implies that Γ_i acts quasi-transitively on G_i for $i = 1, 2$. It remains to show that the identification length is finite. Note that every vertex lies in only finitely many adhesion sets as all those adhesion sets lie in a common Γ -orbit and as G is locally finite. This directly implies that the identification length is finite. \square

6 Accessible graphs

Let G be a connected quasi-transitive locally finite graph with more than one end and let Γ act quasi-transitively on G . We say that G *splits (non-trivially)* into connected quasi-transitive locally finite graphs G_1, G_2 if it is a non-trivial tree amalgamation $G = G_1 * G_2$ of finite adhesion and if the tree-decomposition defined by $G_1 * G_2$ (as in Remark 5.1) is basic with respect to Γ . Note that the stabilizer Γ_i in Γ of G_i acts quasi-transitively on G_i by Proposition 4.5. Now if one of the *factors* G_1 or G_2 also has more than one end, we can split it with respect to Γ_i , too. We can continue this for every factor and call this a *process of splittings*. Note that it is important in a process of splittings to use the group action of the stabiliser of the factor in order to split the factor. If we eventually end up with factors that are either finite or have at most one end, i. e. if the process of splittings terminates, we call the (multi-)set of these factors a *terminal factorisation* of G . (Also, if G is one-ended, we say it is a *terminal factorisation* of itself.) We call G *accessible* if it has a terminal factorisation.

Remark 6.1. Let G be an accessible connected quasi-transitive locally finite graph. Then there are connected quasi-transitive locally finite graphs $G_1, \dots, G_n, H_1, \dots, H_{n-1}$ with $G = H_{n-1}$ and trees T_1, \dots, T_{n-1} such that the following hold:

- (i) every G_i has at most one end;
- (ii) for every $i \leq n - 1$, the graph H_i is a tree amalgamation $H *_{T_i} H'$ with respect to group actions of finite adhesion, where

$$H, H' \in \{G_j \mid 1 \leq j \leq n\} \cup \{H_j \mid 1 \leq j < i\}.$$

Remark 6.2. Let \mathcal{G}_0 be the class of all connected quasi-transitive locally finite graphs with at most one end. For $i > 0$, let \mathcal{G}_i be the class obtained by tree amalgamations of finite adhesion of elements in $\bigcup_{j < i} \mathcal{G}_j$. Set $\mathcal{G} := \bigcup_{i \in \mathbb{N}} \mathcal{G}_i$. Then \mathcal{G} is the class of all accessible connected quasi-transitive locally finite graphs.

A locally finite quasi-transitive graph G is *edge-accessible* if there exists a positive integer k such that any two ends of G can be separated by at most k edges and it is *vertex-accessible* if there exists a positive integer k' such that any two ends of G can be separated by at most k' vertices.

The following result generalizes a graph theoretical characterisation of accessibility of finitely generated groups: Thomassen and Woess [20, Theorem 1.1] proved that a finitely generated group is accessible if and only if it has an edge-accessible locally finite Cayley graph.

Theorem 6.3. *Let G be a connected locally finite quasi-transitive graph. Then the following statements are equivalent.*

- (1) G is accessible.
- (2) G is edge-accessible.
- (3) G is vertex-accessible.

Before we prove Theorem 6.3, we need another result. Recall that a tree-decomposition efficiently distinguishes two ends if there is an adhesion set $V_{t_1} \cap V_{t_2}$ separating them such that no set of smaller size than $V_{t_1} \cap V_{t_2}$ separates them.

Theorem 6.4. *Let G be an vertex-accessible connected locally finite graph and let Γ be a group acting quasi-transitively on G . Then there exists a Γ -invariant tree-decomposition (T, \mathcal{V}) of G of finite adhesion such that (T, \mathcal{V}) distinguishes all ends of G efficiently and such that there are only finitely many Γ -orbits on $E(T)$.*

Proof. Since G is vertex-accessible, there exists $k \in \mathbb{N}$ such that every two ends can be separated by at most k vertices. By Theorem 3.2 we find a Γ -invariant tree-decomposition (T, \mathcal{V}) of G of adhesion at most k that distinguishes all ends efficiently.

For every adhesion set $V_t \cap V_{t'}$ that does not separate any two ends efficiently, we contract the edge tt' in T and assign the vertex set $V_t \cup V_{t'}$ to the new node. It is easy to check that the resulting pair (T', \mathcal{V}') is again a tree-decomposition. It only has adhesion sets that distinguish ends efficiently. Note that Γ still acts on (T', \mathcal{V}') as the set of adhesion sets that do not separate ends efficiently is Γ -invariant. A result of Thomassen and Woess [20, Proposition 4.2] says that there are only finitely many vertex sets S of size at most k containing a fixed vertex such that for two components C_1, C_2 of $G - S$ every vertex of S has a neighbour in C_1 and in C_2 . It follows that there are only finitely many orbits of adhesion sets that separate ends efficiently. This proves the assertion. \square

Proof of Theorem 6.3. The equivalence of (2) and (3) is trivial.

To prove that (3) implies (1), let G be vertex-accessible and let Γ be a group acting on G with only finitely many orbits. By Theorem 6.4 we find a Γ -invariant tree-decomposition (T, \mathcal{V}) of G of finite adhesion such that (T, \mathcal{V}) distinguishes all ends of G efficiently and such that there are only finitely many Γ -orbits on $E(T)$. By Proposition 4.1, we may assume that all adhesion sets are connected. We prove the assertion by induction on the number of Γ -orbits of adhesion sets of (T, \mathcal{V}) . Let $tt' \in E(T)$. For every edge $t_1t_2 \in E(T)$ that does not lie in the same Γ -orbits as tt' , we contract the edge t_1t_2 in T and

assign the vertex set $V_{t_1} \cup V_{t_2}$ to the new node. Let T' be the resulting tree and $\mathcal{V}' = \{V_s \mid s \in V(T')\}$. It is easy to verify that (T', \mathcal{V}') is a tree-decomposition. The only edges of T' are those that have their origin in the Γ -orbit of the edge $tt' \in E(T)$ and Γ still acts on (T', \mathcal{V}') such that (T', \mathcal{V}') is a connected basic tree-decomposition of G with only connected adhesion sets. Lemma 5.4 implies that G is a non-trivial tree amalgamation $G_1 *_{T'} G_2$ with respect to group actions, where the graphs G_1 and G_2 are induced by the parts of (T', \mathcal{V}') . The tree-decomposition (T, \mathcal{V}) induces a tree-decomposition (T_W, \mathcal{W}) on the parts W of (T', \mathcal{V}') and there are less Γ_W -orbits on the adhesion sets of (T_W, \mathcal{W}) than Γ -orbits on the adhesion sets of (T, \mathcal{V}) . Thus, we can apply induction. This shows (1).

To prove that (1) implies (3), we will use the graph classes \mathcal{G}_i and \mathcal{G} as defined in 6.2 and show inductively that every \mathcal{G}_i contains only vertex-accessible connected quasi-transitive locally finite graphs. This is obviously true for \mathcal{G}_0 . Let $G \in \mathcal{G}_i$ for $i > 0$. Then there are $G_1, G_2 \in \bigcup_{j < i} \mathcal{G}_j$ such that G is a tree amalgamation $G_1 *_T G_2$ of finite adhesion. By induction, we may assume that G_1 and G_2 are edge-accessible and quasi-transitive. Note that quasi-transitivity of G follows from Lemma 5.3 since G_1 and G_2 are quasi-transitive. For $i = 1, 2$, let k_i be a positive number such that any two ends of G_i can be separated by at most k_i many vertices. Let (T, \mathcal{V}) be the tree-decomposition we obtain from the tree amalgamation $G_1 *_T G_2$ according to Remark 5.1. Let k be the maximum of k_1, k_2 and the size of adhesion sets of $G_1 *_T G_2$.

Let Q, R be two rays in different ends ω_Q, ω_R of G , respectively. If there is some adhesion set $V_t \cap V_{t'}$ such that Q and R have tails that are separated by $V_t \cap V_{t'}$, then the ends they lie in must be separated by that adhesion set as well. Hence, they are separable by a separator of order at most k . So we may assume that, eventually, they lie on the same side of each separator. By Proposition 4.8 (i) every end of G is captured either by an end or by a node of T . Thus and since no separator separates any tails of Q and R , their ends are captured by the same node or end of T . By Proposition 4.8 (iii) an end of T captures a unique end of G . Thus, ω_Q and ω_R are captured by the same node of T . By Proposition 4.8 (iv) every end of G that is captured by a node $t \in V(T)$ corresponds to a uniquely determined end of $G[V_t]$. These ends can be separated by a separator S in $G[V_t]$ of order at most k by assumption. However, S need not be a separator of G that separates those ends. Still, it is possible to enlarge S to a separator of G that separates ω_Q and ω_R and still has bounded size: every vertex of S has distance at most K , the maximum diameter of the adhesion sets measured in G_1 and in G_2 to only finitely many adhesion sets that are contained in V_t as G is locally finite; so we can add all these adhesion sets to S and obtain a set S' . As G is quasi-transitive, the size of S' only depends on k , the number of orbits of vertices of G , the maximum number of adhesion sets in V_t that have distance at most K to a common vertex and the size of any adhesion set of (T, \mathcal{V}) , in particular, it is bounded by some $\ell \in \mathbb{N}$ and it is independent of the chosen ends. If we show that S' separates ω_Q and ω_R , then it follows immediately that G is edge-accessible.

Let $P = \dots, x_{-1}, x_0, x_1, \dots$ be a double ray with its tail x_0, x_1, \dots in ω_Q and its tail x_0, x_{-1}, \dots in ω_R . Since both ends ω_Q and ω_R are captured by V_t , there are infinitely many x_i with $i > 0$ that lie in V_t and infinitely many x_i with $i < 0$ that lie in V_t . Let us assume $x_0 \in V_t$. Whenever the ray $P^+ := x_0 x_1 \dots$ leaves V_t through an adhesion set $V_t \cap V_{t'}$, it must reenter V_t and this must happen

through the same adhesion set. Since S is finite and separates ω_Q and ω_R , there are $i_1, i_2 \in \mathbb{Z}$ such that no $x_i \in V_t$ with $i \geq i_1$ is separated in $G[V_t]$ by S from ω_Q and no $x_i \in V_t$ with $i \leq i_2$ is separated in $G[V_t]$ by S from ω_R . Then there must be some path x_i, \dots, x_j with $j \geq i+1$ and whose inner vertices lie outside of V_t such that x_j is not separated by S from ω_Q and x_i is not separated by S from ω_R . Thus, the shortest x_i - x_j path in $G[V_t]$ meets S . As x_i and x_j lie in a common adhesion set, we conclude that this lies in S' . Thus, S' separates ω_Q from ω_R in G . This shows (3). \square

In the proof of the implication (3) to (1) of Theorem 6.3 we chose a specific way to split the factors. (It was based on a Γ -invariant tree-decomposition of G .) We do not know if we can split arbitrary in each step and still have to end in a terminal factorisation. But we conjecture that this is true.

Conjecture 6.5. *Let G be an accessible connected quasi-transitive locally finite graph. Every process of splittings must end after finitely many steps.*

Accessibility of finitely generated groups received lots of attention after Wall [21] conjectured that all finitely generated groups are accessible and among the main results in this area are Dunwoody's results that Wall's conjecture is false in general [8] but true for (almost) finitely presented groups [7]. In the case of quasi-transitive locally finite graphs, the investigation focused on edge-accessible graphs, see [11, 17, 20]. However, Theorem 6.3 enables us to carry over these results to accessible graphs.

7 Applications

7.1 Stallings' theorem

In this section we will discuss how to obtain Stallings' theorem from our results.

Let Γ be a finitely generated group with infinitely many ends and let G be a locally finite Cayley graph of Γ . Then G has infinitely many ends, too. By Theorem 5.5, G is a non-trivial tree amalgamation $G_1 *_T G_2$ of finite adhesion. Since it has finite adhesion and Γ acts regularly⁴ on G , the stabiliser in Γ of an edge of T , which is just the stabiliser in Γ of the corresponding adhesion set, is finite. Hence, Bass-Serre theory leads to Stallings' theorem.

Theorem 7.1. [18] *If a finitely generated group has more than one end, then it is either a free product with amalgamation over a finite subgroup or an HNN-extension over a finite subgroup.* \square

Note that tree amalgamations of Type 1 respecting the actions of groups acting on the factors lead to free products with amalgamation and tree amalgamations of a graph with itself with respect to the action of a group leads to an HNN-extension.

7.2 Graphs without thick ends

Let us apply our main results to connected quasi-transitive locally finite graphs that have only thin ends. First, we want to see that such graphs are edge-

⁴i. e. for every two $u, v \in V(G)$ there is a unique element of Γ mapping u to v

accessible, so we can apply Theorem 6.4 and look at terminal factorisations of them. But before we go into the proof, we need some definitions.

Let G and H be graphs. A map $\varphi: V(G) \rightarrow V(H)$ is a (γ, c) -*quasi-isometry* if there are constants $\gamma \geq 1, c \geq 0$ such that

$$\gamma^{-1}d_G(x, y) - c \leq d_H(\varphi(x), \varphi(y)) \leq \gamma d_G(x, y) + c$$

for all $x, y \in V(G)$ and such that $\sup\{d_H(x, \varphi(V(G))) \mid x \in V(H)\} \leq c$. We then say that G is *quasi-isometric* to H .

Krön and Möller [13, Theorem 5.5] showed that a connected quasi-transitive locally finite graph has only thin ends if and only if it is quasi-isometric to a tree. Trees are obviously edge-accessible and it follows from the definition of edge-accessibility that the class of edge-accessible quasi-transitive locally finite graphs is invariant under quasi-isometries. Thus, we have verified the following.

Proposition 7.2. *Every connected locally finite quasi-transitive graph that has only thin ends is edge-accessible.* \square

We mention that Thomassen and Woess [20, Theorem 5.3] showed Proposition 7.2 for transitive graphs directly with a nice graph theoretical argument. It is not too hard to modify their argument in such a way that the proof works for quasi-transitive graphs as well.

Another result we need for our investigation here is due to Thomassen.

Proposition 7.3. [19, Proposition 5.6.] *If G is an infinite connected quasi-transitive locally finite graph with only one end, then the end is thick.* \square

Recently, Carmesin et al. [3, Theorem 5.1] extended Proposition 7.3 to graphs that need not be locally finite.

Now we are able to give a new characterisation of connected quasi-transitive locally finite graphs with only thin ends.

Theorem 7.4. *A connected quasi-transitive locally finite graph has only thin ends if and only if it has a terminal factorisation of only finite graphs.*

Proof. Let G be a connected quasi-transitive locally finite graph. First, let us assume that every end of G is thin. By Theorem 7.2, G is edge-accessible. So Theorem 6.4 implies that G is accessible and hence has a terminal factorisation. All the factors of that terminal factorisation have at most one end. Since they are quasi-transitive by Proposition 4.5, they cannot have one end due to Proposition 7.3. So they are locally finite graphs without ends, which implies that they are finite graphs.

For the other direction, we follow the steps to factorise G , factorise each of its factors and so on until we end up with a terminal factorisation. Note that by Proposition 4.8 (ii) every thick end of G is captured by nodes of the involved basic tree-decompositions. So if G had a thick end, then one of the factors of the terminal factorisation must have a thick end, which is impossible as these factors are finite by assumption. Thus, all ends of G are thin. \square

Note that there are several characterisations of (quasi-transitive or Cayley) graphs that are quasi-isometric to trees, see e.g. Antolín [1] and Krön and Möller [13]. We enlarged their list of characterisations by our theorem.

A natural class of quasi-transitive graphs are Cayley graphs. So our theorems apply in particular for such graphs and we obtain as a corollary of Theorem 7.4 a result for virtually free groups. A group Γ is *virtually free* if it contains a free subgroup of finite index.

Woess [22] showed that a finitely generated group is virtually free if and only if every end of any of its locally finite Cayley graphs is thin. Thus we directly obtain the following corollary.

Corollary 7.5. *A finitely generated group is virtually free if and only if any of its locally finite Cayley graphs has a terminal factorisation of only finite graphs.* \square

In [10] the interplay between tree amalgamations and quasi-isometries is investigated further and the results of this section are extended to graphs other than trees in two ways. First, it is shown that the quasi-isometry type of (iterated) tree amalgamations only depend on the quasi-isometry types of the infinite factors. Then, in the case of accessible infinitely-ended graphs, it is shown that the quasi-isometry types of the graphs determine the quasi-isometry types of the infinite factors in any of its terminal factorisations.

7.3 Planar graphs

Mohar, see [14], raised the question whether tree amalgamations are powerful enough to characterise planar transitive locally finite graphs in terms of finite or one-ended locally finite planar transitive graphs. The aim of this section is to answer his question in the affirmative in case of planar quasi-transitive graphs.

Dunwoody [9] proved that planar quasi-transitive locally finite graphs are edge-accessible, see also [12]. This allows us to apply Theorem 6.3 to these graphs. We directly obtain the following result.

Theorem 7.6. *For every planar connected quasi-transitive locally finite graph G there are finitely many planar connected quasi-transitive locally finite graphs G_1, \dots, G_n with at most one end such that G can be obtained by finitely many (iterated) tree amalgamations of G_1, \dots, G_n .* \square

7.4 Hyperbolic graphs

For our last application, we look at hyperbolic graphs. The aim is to give a characterisation of quasi-transitive locally finite hyperbolic graphs in terms of their terminal factorisations. But before we state the main theorem of this section, we need some definitions and preliminary results.

Let $\delta \geq 0$ and let G be a graph. Then G is δ -*hyperbolic* if for all vertices $x_1, x_2, x_3 \in V(G)$ and all geodesics $P_{i,j}$ between x_i and x_j , every vertex of the path $P_{1,2}$ has distance at most δ to a vertex on $P_{2,3} \cup P_{1,3}$. We call G *hyperbolic* if it is δ -hyperbolic for some $\delta \geq 0$.

Let $\gamma \geq 1$ and $c \geq 0$. A finite walk is called a (γ, c) -*quasi-geodesic* if it is the image of a (γ, c) -quasi-isometry of some geodesic with the same end vertices.

Lemma 7.7. *Let G_1, G_2 be connected locally finite graphs and let $G = G_1 * G_2$ be a tree amalgamation of G_1 and G_2 such that the adhesion sets in G_1 have bounded diameter in G_1 . Then there are some $\gamma \geq 1$, $c \geq 0$ such that every geodesic in G_1 is a (γ, c) -quasi-geodesic in G .*

Proof. Let γ' be the maximum distance in G_1 between vertices of an adhesion set. If $\gamma' = 1$, then it is straightforward to see that every geodesic in G_1 is a geodesic in G . So we may assume $\gamma' \geq 2$. Let P_1 be a geodesic in G_1 and let P be a geodesic in G with the same end vertices as P_1 . Whenever P contains a vertex outside of G_1 , it must have left G_1 through an adhesion set and reentered through the same adhesion set. (Note that we consider the tree-decomposition defined by the tree amalgamation $G_1 * G_2$ as in Remark 5.1.) We replace every maximal subpath of P all whose inner vertices lie outside of G_1 by a path in G_1 with the same end vertices. As these end vertices lie in a common adhesion set, the length of the replacement path is at most γ' and it replaces a path of length at least 2. We end up with a walk in G_1 that has the same end vertices as P and that is at most $\gamma'/2$ times as long as P . Since P_1 is shorter than that walk, its length is at most $\gamma'/2$ times as long as P . As the same applies for every subpath of P_1 , we conclude that P_1 is a $(\gamma'/2, 0)$ -quasi-geodesic in G . \square

The reason that we are interested in quasi-geodesics is because they still lie relatively close to geodesics in hyperbolic graphs as the following lemma shows.

Lemma 7.8. [4, Théorème 3.1.4] *Let G be a locally finite δ -hyperbolic graph. For all $\gamma \geq 1$ and $c \geq 0$, there is a constant $\kappa = \kappa(\delta, \gamma, c)$ such that for every two vertices x, y of G every (γ, c) -quasi-geodesic between them lies in a κ -neighbourhood around every geodesic between x and y and vice versa.* \square

Lemma 7.9. *Let (T, \mathcal{V}) be a connected basic tree-decomposition of a connected locally finite graph G such that every part induces a hyperbolic graph. Let $x, y \in V_t$ for some $t \in V(T)$. Then there exists $\lambda = \lambda(\delta, \gamma)$ such that every geodesic in G between x and y lies in a λ -neighbourhood of every geodesic in $G[V_t]$ between x and y and vice versa, where γ is the diameter of any adhesion set.*

Proof. Let P, P' be a geodesic in G , in $G[V_t]$ between x and y , respectively. We modify P the same way we did it in the proof of Lemma 7.7: we replace every maximal subpath of P all whose inner vertices lie outside of $G[V_t]$ by a shortest path in $G[V_t]$ with the same end vertices. This is possible as whenever P contains a vertex outside of $G[V_t]$ it must have left $G[V_t]$ through an adhesion set and reentered through the same adhesion set. Let P'' be the walk obtained from P after these replacements. If $\gamma = 1$, then $P = P''$ and if $\gamma \neq 1$, then the length of P'' is at most $\gamma/2$ times the length of P as paths of length at least 2 got replaced by paths of length at most γ . Set $\gamma' := \max\{1, \gamma/2\}$. Now let a, b be vertices of P'' . Then there are vertices a', b' on $P \cap P''$ of distance at most $\gamma/2$ to a , to b , respectively. Hence, we have

$$d_{P''}(a, b) \leq \gamma + d_{P''}(a', b') \leq \gamma + \gamma' d_G(a', b') \leq (1 + \gamma')\gamma + \gamma' d_G(a, b).$$

So P'' is $(\gamma', (1 + \gamma')\gamma)$ -quasi-geodesic. Applying Lemma 7.8, we find κ depending only on δ and γ such that P'' lies in a κ -neighbourhood of P' and vice versa. Since P lies in a $\gamma/2$ -neighbourhood of P'' and vice versa, we have shown the assertion. \square

The following theorem is our main result for quasi-transitive locally finite hyperbolic graphs: it shows that tree amalgamations behave well with respect to hyperbolicity.

Theorem 7.10. *Let G_1 and G_2 be connected locally finite graphs and let G be a tree amalgamation $G_1 * G_2$ such that the adhesion sets in G_i have bounded diameter in G_i for $i = 1, 2$. Then G is hyperbolic if and only if G_1 and G_2 are hyperbolic.*

Proof. First, let G be δ -hyperbolic. Let $x_1, x_2, x_3 \in V(G_1)$ and let P_{ij} be a geodesic in G_1 and P'_{ij} be a geodesic in G between x_i and x_j for all $i \neq j$. Let $x \in P_{12}$. By Lemma 7.7 and its proof, P_{12} is $(\gamma, 0)$ -quasi-geodesic for $\gamma = \max\{1, \beta/2\}$, where β is the maximum distance in G_1 between two vertices in a common adhesion set in G_1 . By Lemma 7.8, there is some $x' \in P'_{12}$ of distance at most κ to x for some $\kappa \geq 0$. Since G is δ -hyperbolic, we find $y' \in P'_{13} \cup P'_{23}$ of distance at most δ to x' . Again by Lemma 7.8, we find $y \in P_{13} \cup P_{23}$ of distance at most κ to y' . Hence, we have $d_G(x, y) \leq 2\kappa + \delta$. Lemma 7.7 implies $d_{G_1}(x, y) \leq \gamma(2\kappa + \delta)$. Thus, G_1 is $\gamma(2\kappa + \delta)$ -hyperbolic. Analogously, G_2 is $\gamma(2\kappa + \delta)$ -hyperbolic.

Now let G_1 and G_2 be hyperbolic. Then there is some $\delta \geq 0$ such that G_1 and G_2 are δ -hyperbolic. Let γ_i be the maximum distance between vertices in a common adhesion set in G_i for $i = 1, 2$ and let $\gamma := \max\{\gamma_1, \gamma_2\}$. We consider the canonical tree-decomposition (T, \mathcal{V}) as discussed in Remark 5.1. Note that the parts of (T, \mathcal{V}) induce graphs that are isomorphic to either G_1 or G_2 , so they are hyperbolic. Let $x_1, x_2, x_3 \in V(G)$ and let P_{ij} be a geodesic between x_i and x_j . Let $t_1, t_2 \in V(T)$ of minimum distance to each other with $x_i \in V_{t_i}$ for $i = 1, 2$ and let T_{12} be the t_1 - t_2 path in T . Note that every node on T_{12} and every adhesion set $V_t \cap V_{t'}$ with $tt' \in E(T_{12})$ contains a vertex of P_{12} and also of $P_{13} \cup P_{23}$.

Let $x \in P_{12}$. Let $t' \in V(T)$ closest to T_{12} with $x \in V_{t'}$ and let $t \in T_{12}$ closest to t' . We say P_{12} passes through V_t in parallel to P_{13} either if $t = t_1$ and P_{13} contains a vertex of the adhesion set $V_t \cap V_{t'_2}$, where t'_2 is the neighbour of t on T_{12} , or if t is neither t_1 nor t_2 and both $V_t \cap V_{t'_1}$ and $V_t \cap V_{t'_2}$ contain vertices of P_{13} , where t'_i is the neighbour of t on T_{12} closest to t_i for $i = 1, 2$. Analogously, P_{12} passes through V_t in parallel to P_{23} either if $t = t_2$ and P_{23} contains a vertex of the adhesion set $V_t \cap V_{t'_1}$, where t'_1 is the neighbour of t on T_{12} , or if t is neither t_1 nor t_2 and both $V_t \cap V_{t'_1}$ and $V_t \cap V_{t'_2}$ contain vertices of P_{23} , where t'_i is the neighbour of t on T_{12} closest to t_i for $i = 1, 2$.

First, let us assume that P_{12} passes through V_t in parallel to P_{13} . If $t = t_1$, let $u_1 := v_1 := x_1$, let u_2 be the last vertex on P_{12} in $V_t \cap V_{t'_2}$, and let v_2 be a vertex on P_{13} in $V_t \cap V_{t'_2}$. If $t \neq t_1$, let u_1 be the first vertex on P_{12} in $V_t \cap V_{t'_1}$ and u_2 be the last vertex on P_{12} in $V_t \cap V_{t'_2}$ and let v_1, v_2 be on P_{13} in $V_t \cap V_{t'_1}$, in $V_t \cap V_{t'_2}$, respectively. Note that by the choice of u_1 and u_2 , the vertex x lies between u_1 and u_2 on P_{12} . Let P be a geodesic in V_t between u_1 and u_2 and let P' be a geodesic in V_t between v_1 and v_2 . Let Q_i be a geodesic in V_t between u_i and v_i for $i = 1, 2$. Looking at a v_1, v_2, u_2 , we conclude by δ -hyperbolicity that any geodesic between v_1 and v_2 lies in a δ -neighbourhood of $Q_2 \cup P'$, so P lies in a 2δ -neighbourhood of $Q_1 \cup P' \cup Q_2$. As the lengths of Q_1 and of Q_2 are bounded by γ , we conclude that P lies in a $(2\delta + \gamma)$ -neighbourhood of P' . Lemma 7.9 implies the existence of some λ such that P contains a vertex y_1 of distance at most λ from x . We just showed that P' contains a vertex y_2 of distance at most $2\delta + \gamma$ from y_1 and Lemma 7.9 shows the existence of a vertex y_3 on P_{23} with $d(x, y_3) \leq 2\lambda + 2\delta + \gamma$.

Analogously, we conclude in the case that P_{12} passes through V_t in parallel

to P_{23} that P_{23} contains a vertex of distance at most $d(x, y_3) \leq 2\lambda + 2\delta + \gamma$ from x .

Let us now assume that P_{12} passes through V_t neither in parallel to P_{13} nor in parallel to P_{23} . If $t = t_1$, let $u_1 := v_1 := x_1$. If $t \neq t_1$, let t'_1 be the neighbour of t on T_{12} closest to t_1 and let u_1 be the first vertex on P_{12} in $V_t \cap V_{t'_1}$ and let v_1 be a vertex on P_{13} in $V_t \cap V_{t'_1}$. If $t = t_2$, let $u_2 := w_2 := x_1$. If $t \neq t_2$, let t'_2 be the neighbour of t on T_{12} closest to t_2 and let u_2 be the last vertex on P_{12} in $V_t \cap V_{t'_2}$ and let w_2 be a vertex on P_{23} in $V_t \cap V_{t'_2}$. Let $t_3 \in V(T)$ of minimum distance to t such that $x_3 \in V_{t_3}$. If $t = t_3$, let $w_1 := v_2 := x_3$. If $t \neq t_3$, let t'_3 be the neighbour of t in T closest to t_3 and let w_1 be a vertex on P_{13} in $V_t \cap V_{t'_3}$ and let v_2 be a vertex on P_{23} in $V_t \cap V_{t'_3}$. (Note that both P_{13} and P_{23} must pass through the adhesion set $V_t \cap V_{t'_3}$ according to the definition of a tree-decomposition.)

We consider a couple of geodesics in $G[V_t]$: let P_u, P_v, P_w be a geodesic in $G[V_t]$ between u_1, u_2 , between v_1, v_2 , between w_1, w_2 , respectively and let P_{uv}, P_{vw}, P_{uw} be a geodesic in $G[V_t]$ between u_1 and v_1 , between v_2 and w_1 , between u_2 and w_2 , respectively. Similar to the case that P_{12} passes through V_t in parallel to P_{13} we conclude that P_u lies in a 4δ -neighbourhood of $P_{uv} \cup P_v \cup P_{vw} \cup P_w \cup P_{uw}$ and hence in a $(4\delta + \gamma)$ -neighbourhood of $P_v \cup P_w$.

Let λ be the value obtained in Lemma 7.9. Then there is a vertex y_1 on P_u of distance at most λ from x . As we just showed, we find y_2 on either P_u or P_w with $d(y_1, y_2) \leq 4\delta + \gamma$ and Lemma 7.9 then implies the existence of a vertex y_3 on either P_{13} or P_{23} with $d(y_2, y_3) \leq \lambda$. So we have $d(x, y_3) \leq 2\lambda + 4\delta + \gamma$. This proves that G is $(2\lambda + 4\delta + \gamma)$ -hyperbolic. \square

As a corollary of Theorem 7.10, we obtain a characterisation of quasi-transitive locally finite hyperbolic graphs in terms of their terminal factorisations.

Corollary 7.11. *A connected quasi-transitive locally finite graph is hyperbolic if and only if it admits a terminal factorisation such that all its factors are connected quasi-transitive locally finite hyperbolic graphs with at most one end.*

Proof. Let G be a connected quasi-transitive locally finite graph. If G is one-ended, then it is a terminal factorisation of itself and the assertion holds trivially. So let us assume that G has more than one end.

First, let us assume that G is hyperbolic. By [11, Theorem 4.3], it is an edge-accessible graph. Thus it is accessible and has a terminal factorisation by Theorem 6.3. So there are connected quasi-transitive locally finite graphs $G_1, \dots, G_n, H_1, \dots, H_{n-1}$ with $G = H_{n-1}$ such that each G_i has at most one end and for every $i \leq n-1$, the graph H_i is a tree amalgamation $H * H'$ of finite adhesion, where

$$H, H' \in \{G_j \mid 1 \leq j \leq n\} \cup \{H_j \mid 1 \leq j < i\}.$$

(We may assume that all G_i are indeed needed at some point during these tree amalgamations.) By repeated application of Lemma 7.10, each H_i , and thus each G_i is hyperbolic.

Conversely, if G has a terminal factorisation into connected finite or connected quasi-transitive locally finite hyperbolic one-ended graphs, then each of the previous factors we considered for obtaining the terminal factorisation are hyperbolic by Lemma 7.10. In particular, G is hyperbolic. \square

References

- [1] Y. Antolín. On cayley graphs of virtually free groups. *Groups–Complexity–Cryptology*, 3(2):301–327, 2011.
- [2] J. Carmesin, M. Hamann, and B. Miraftab. Canonical tree-decompositions of infinite graphs. in preparation.
- [3] J. Carmesin, F. Lehner, and R.G. Möller. On tree-decompositions of one-ended graphs. arXiv:1706.08330.
- [4] M. Coornaert, T. Delzant, and A. Papadopoulos. *Géométrie et théorie des groupes. Les groupes hyperboliques de Gromov*, volume 1441 of *Lecture notes in Math.* Springer-Verlag, 1990.
- [5] R. Diestel. *Graph Theory*. Springer, 4th edition, 2010.
- [6] M.J. Dunwoody. Cutting up graphs. *Combinatorica*, 2:15–23, 1982.
- [7] M.J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81(3):449–457, 1985.
- [8] M.J. Dunwoody. An inaccessible group. In G.A. Niblo and M.A. Roller, editors, *Geometric Group Theory*, volume 181 of *L.M.S. Lecture Note Ser.*, pages 75–78. Cambridge University Press, 1992.
- [9] M.J. Dunwoody. Planar graphs and covers. preprint, 2007.
- [10] M. Hamann. Tree amalgamations and quasi-isometries. arXiv:1812.04987.
- [11] M. Hamann. Accessibility in transitive graphs. *Combinatorica*, 38:847–859, 2018.
- [12] M. Hamann. Planar transitive graphs. *Electronic J. Combin.*, 25:Paper 4.8, 2018.
- [13] B. Krön and R.G. Möller. Quasi-isometries between graphs and trees. *Journal of Combinatorial Theory, Series B*, 98(5):994–1013, 2008.
- [14] B. Mohar. Tree amalgamation of graphs and tessellations of the Cantor sphere. *Journal of Combinatorial Theory, Series B*, 96(5):740–753, 2006.
- [15] R. Möller. Ends of graphs. *Math. Proc. Cambridge Phil. Soc.*, 111:255–266, 1992.
- [16] R. Möller. Ends of graphs II. *Math. Proc. Cambridge Phil. Soc.*, 111:455–460, 1992.
- [17] R.G. Möller. Accessibility and ends of graphs. *J. Combin. Theory (Series B)*, 66(2):303–309, 1996.
- [18] J.R. Stallings. *Group theory and three-dimensional manifolds*, volume 4. Yale Univ, Press, New Haven, CN, 1971.
- [19] C. Thomassen. The Hadwiger number of infinite vertex-transitive graphs. *Combinatorica*, 12(4):481–491, 1992.

- [20] C. Thomassen and W. Woess. Vertex-transitive graphs and accessibility. *J. Combin. Theory (Series B)*, 58:248–268, 1993.
- [21] C.T.C. Wall. Pairs of relative cohomological dimension one. *J. Pure Appl. Algebra*, 1:141–154, 1971.
- [22] W. Woess. Graphs and Groups with Tree-like Properties. *J. Combin. Theory Ser. B*, 47(3):361–371, 1989.