

VERTEX CUTS

M.J. DUNWOODY AND B. KRÖN

Version 30.4.2009

ABSTRACT. We generalise structure tree theory, which is based on removing finitely many edges, to removing finitely many vertices. This gives a significant generalization of Tutte’s tree decomposition of 2-connected graphs into 3-connected blocks. For a finite graph there is a structure tree that contains information about k -connectivity for any k . The theory can also be applied to infinite graphs that have more than one vertex end, i.e. ends that can be separated by removing a finite number of vertices. This gives a generalization of Stallings’ structure theorem for groups with more than one end.

1. INTRODUCTION

A connected simple graph X is said to be n -connected if for every pair u, v of distinct vertices there are n paths joining u to v such that if $p \in VX$ and $p \neq u, p \neq v$ then p lies on at most one of the paths. If X is not 2-connected then it has *cut-points*, i.e. vertices whose removal disconnects the graph. If this happens, and X has no disconnecting edges, then X decomposes into a collection of maximal 2-connected subgraphs, or *2-blocks*. Any two 2-blocks intersect in at most one vertex and this vertex will be a cut-point. Every edge of X lies in exactly one 2-block. Associated with this decomposition is a *structure tree* T in which $VT = \mathcal{B} \cup \mathcal{S}$, where \mathcal{S} is the set of cut-points, \mathcal{B} is the set of 2-blocks and there is an edge joining $b \in \mathcal{B}$ with $s \in \mathcal{S}$ if and only if $s \in b$. If G is a group of automorphisms of X , then there is an induced action of G on T . If X is a finite graph then T will be a finite tree and any action on a finite tree is trivial, i.e. there is a vertex which is fixed by G . This is illustrated in Figure 1.




FIGURE 1. One-connected graph and structure tree

The number next to a vertex of T indicates the order of the subgroup of the automorphism group fixing that vertex. Note the 2-colouring of T , in which white vertices are cut points and black vertices are blocks.

There is a similar decomposition if X is 2-connected but not 3-connected. This was described by Tutte [25] if X is finite and Droms Servatius and Servatius [3] if X is infinite and locally finite. A somewhat different account is given in [7]. The decomposition gives a structure tree T , which again has a 2-colouring $VT = \mathcal{B} \cup \mathcal{S}$. In this case each vertex $s \in \mathcal{S}$ corresponds to a 2-separator, i.e. a pair of vertices whose removal disconnects the graph, and for each maximal 3-connected subgraph or 3-block there is a vertex $b \in \mathcal{B}$. In this case the situation is more complicated than for cut-points as there are also black vertices which do not correspond to 3-blocks. An example is given in Figure 2. Here the black central vertex of T corresponds to a 4-cycle one edge of which, shown dashed, is what Tutte calls an ideal edge. This edge is not in the original graph, and joins the vertices of a 2-separator. In the structure tree that he constructs every black vertex corresponds to a 3-block or a cycle in which some edges may be ideal edges. Note that we could add ideal edges so that all the blocks were complete graphs as in Figure 3.




FIGURE 2. Decomposition of a 2-connected graph with one ideal edge.

In this paper we show that if a graph X has a finite set of vertices whose removal produces at least two components that are large in some sense and G is the automorphism group of X then there is a G -tree (or structure tree) T with a bipartition $(\mathcal{S}, \mathcal{B})$ of the set of vertices $VT = \mathcal{S} \cup \mathcal{B}$ so that the vertices in \mathcal{S} correspond to finite separating sets.




FIGURE 3. Decomposition of a 2-connected graph with all ideal edges.

Such a structure tree had been known to exist in the case when X is an infinite graph and X can be disconnected into two infinite components by removing finitely many edges (see [2, 4]).

A *ray* is a sequence of distinct vertices v_0, v_1, \dots such that v_i and v_{i+1} are adjacent for each i . Let $E \subset VX$. The *coboundary* δE is the set of edges that have one vertex in E and one in $VX \setminus E$. If δE is finite, then E is called an *edge cut*. If R is a ray, then all the terms from a certain point onwards will either lie in E or in $VX \setminus E$. We say that E separates rays R_1, R_2 if one of the rays eventually belongs to E and the other eventually belongs to $VX \setminus E$. We say two rays belong to the same edge end if they are not separated by any edge cut. It is easy to see that this is an equivalence relation on the set of rays, and so we can take an equivalence class to be an *edge end*.

In [4] it is shown that if a graph has more than one edge end and k is the smallest integer for which there is an edge cut E with $|\delta E| = k$ that separates two ends, then there is a structure tree in which the edges correspond to edge cuts with this property. In [2] there has been a substantial theory developed for edge cuts and edge ends starting from Stallings' Theorem on the structure of groups with infinitely many ends ([21, 23]).

In this paper we are concerned with vertex cuts and vertex ends. We say that $A \subset VX$ is a vertex cut if A is connected, i.e. any two vertices can be joined by a path all of whose vertices are in A , and VX can be partitioned $A \cup NA \cup A^*$, where NA is finite and consists of the vertices which are not in A , but which are adjacent to vertices in A . Note that generally A^* will not be connected. As for edge cuts any ray is eventually in A or in A^* . We say two rays belong to the same vertex end if they are not separated by any vertex cut. A finite set F of vertices is called a *separator* if $VX \setminus F$ has at least two components which contain an end. If E is an edge cut then it is also a vertex cut in which NE is the set of vertices of δE which are not in E . Thus if two rays belong to the same vertex end, then they belong to the same edge end. The converse is true if X is locally finite. However it is easy to construct examples of graphs which are not locally finite in which there are more vertex ends than edge ends. For example if K_∞ is the complete graph on a countably infinite set of vertices and X is the graph consisting of n copies of K_∞ , in which a single vertex from each copy is identified, then X has n vertex ends but only one edge end.

Another example is the Farey Graph as illustrated in Figure 4. This was pointed out to us by Hamish Short. This graph is obtained by taking an ideal triangle in the hyperbolic plane and then taking all translates of this triangle under the group of isometries generated by reflexions in the three sides. One obtains a graph, in which the vertices are the translates of the vertices of the triangle. All of these will lie in the boundary of the plane, which will be a circle in the disc model. The edges of the graph will be the translates of the edges of the triangle. The vertices of any edge will form a 2-separator. In this graph every vertex has infinite valency. The structure tree is easy to see. There will be one orbit of vertices corresponding to the separating edges. The other orbit corresponds to the triangles. For each such triangle there will be three edges of the structure tree joining the vertex corresponding to the triangle to the three vertices corresponding to its boundary edges. The structure tree is essentially the dual graph to the tessellation of the




FIGURE 4. Farey graph with structure tree.

hyperbolic plane. This dual graph is a tree since each edge of the original graph is separating.

In developing our theory we give a set of axioms that it is sufficient for a set of vertex cuts to satisfy in order that a structure tree be constructed. In Figure 2 removing any two of the central four vertices will leave two components. The 12 components thus obtained satisfy the axioms of a *cut system*. The 12 cuts are not *nested* with each other. Thus if one takes two cuts E, F such that NE, NF are the two distinct diagonal pairs, then $E \cap F, E \cap F^*, E^* \cap F, E^* \cap F^*$ are all non-empty intersections. If we restrict to the components obtained by removing two adjacent vertices in the central cycle then we obtain a nested cut system, and the cuts in this system can be regarded as the directed edges of the structure tree.

If X is an infinite graph with more than one vertex end and k is the smallest integer for which there is a vertex cut A such that $|NA| = k$ and k separates two ends then there is a set of such vertex cuts which satisfies the axioms.

One big advantage of using vertex cuts over edge cuts is that we can obtain a structure tree theory that applies to finite graphs, and gives information about the k -connectivity of the graph for any k . For a complete graph the structure tree is trivial, i.e. it only has one vertex. A finite graph X has a non-trivial structure tree if and only if for some integer k , there are k -separators. A k -separator is a set S of k vertices whose removal leaves at least two components C, D such that $C \cup S$ and $D \cup S$ each contain k -inseparable subsets. Here a set Y of vertices is k -inseparable if Y has at least $k + 1$ vertices and no two of the vertices will lie in distinct components of the graph when at most k vertices of X are removed. Let κ be the smallest value of k for which the above occurs. We show that a finite graph contains a unique nested set \mathcal{E} of κ -separators such that if two κ -inseparable subsets are separated by some κ -separator then they are separated by a set in \mathcal{E} . The set \mathcal{E} is invariant under the automorphism group of X and forms the directed edge set of a structure tree for X .

In recent work A. Evangelidou and P. Papasoglu [8] have independently obtained some of the results of this paper. In particular they have obtained a new proof of Stallings' Theorem using vertex cuts.

The authors first met at the conference on Totally Disconnected Groups, Graphs and Geometry at Blaubeuren in May 2007. This project grew from a problem

raised at a problem session at that conference. The authors are very grateful to the organizers for inviting them to the meeting.

2. SYSTEMS OF CUTS AND SEPARATORS

The *boundary* NE of a set of vertices E is the set of vertices in $VX \setminus E$ which are adjacent to E . Set $E^* = VX \setminus (E \cup NE)$. We call E^* the **-complement* of E .

Let E and F be sets of vertices. The intersections $E \cap F$, $E^* \cap F$, $E \cap F^*$ and $E^* \cap F^*$ are called the *corners* of E and F , see Figure 5. The sets $E \cap NF$, $E^* \cap NF$, $F \cap NE$ and $F^* \cap NE$ are called the *links* and $NE \cap NF$ is the *centre*. A link and a corner are said to be *adjacent* if they are adjacent in Figure 5. We say that two links are the links *of* their adjacent corner (or we say they are *its* links), and we say that two corners are the corners *of* their adjacent link (or *its* corners). Two links or two corners are said to be *adjacent* if they are adjacent to the same link or corner, respectively. Otherwise they are called *opposite*.

Let \mathcal{C} be set of non-empty connected sets of vertices with finite boundaries in a connected graph. We call \mathcal{C} a *cut system* if it satisfies the following axioms.

- (A1) If C is in \mathcal{C} then C^* contains an element of \mathcal{C} .
- (A2) If C is in \mathcal{C} then every component of C^* which contains an element of \mathcal{C} is in \mathcal{C} .
- (A3) If C and D are in \mathcal{C} then either a component of $C \cap D$ and a component of $C^* \cap D^*$ are in \mathcal{C} or a component of $C \cap D^*$ and a component of $C^* \cap D$ are in \mathcal{C} .

If \mathcal{C} is a cut system then the boundary of a cut is called a *separator* and we denote the system of separators by \mathcal{S} . Note that in general a given system of separators does not determine a cut system.

A set of vertices E is called *large* if both E and E^* have components which are cuts, otherwise it is called *small*. The **-complement* of a cut is large but it is not necessarily a cut, because it is not necessarily connected.

Axiom (A3) says that if C , D are cuts then there is a pair of large opposite corners of C and D . Or in other words, a large set minus a separator is still large.

Let us consider some examples of cut systems before developing the theory.

Example 2.1. For infinite graphs with more than one (vertex) end, take cuts to be connected sets of vertices with finite boundary which contain a ray (that is, an end) and whose complement also contains a ray. We call these cuts the *vertex end cuts*. Axiom (A3) follows because if C , D , C^* , D^* all contain rays, then so do two opposite corners.

Example 2.2. For infinite graphs with more than one edge end, the separators would naturally be finite sets of edges which separate rays. But separators are by definition sets of vertices. Hence we replace every edge of the original graph by paths of length two. Let M be the set of new vertices, that is, the set of middle vertices of these paths of length two. Then we take cuts to be connected sets of vertices which contain a ray, whose complement also contains a ray and whose boundary is a finite subset of M .

Examples 2.1 and 2.2 can be generalized as follows.

Example 2.3. Let X be a connected graph and let $M \subset VX$ be some set of vertices. Take cuts to be connected sets of vertices which contain a ray, whose complement also contains a ray and whose boundary is a finite subset of M .

Next we consider a cut system which also makes sense in finite graphs. Let k be a positive integer. A subset Y of VX is said to be k -inseparable if it has at least $k + 1$ elements and if for every set $A \subset VX$ with $|NA| \leq k$, either $Y \subset A \cup NA$ or $Y \subset A^* \cup NA$. Examples of k -inseparable subgraphs are the vertex set of a $(k + 1)$ -connected subgraph, or the vertex set of a subgraph which is complete on $k + 1$ vertices. The vertices of a separating edge form a maximal 1-inseparable set.

Example 2.4. Let κ be the smallest positive integer for which there are sets A , Y_1 and Y_2 such that $|NA| = \kappa$, Y_1 and Y_2 are κ -inseparable, $Y_1 \subset A \cup NA$ and $Y_2 \subset A^* \cup NA$. We will show shortly that the connected sets A with this property form a cut system.

Note that there are graphs where this cut system is empty. For instance consider a finite complete graph.

We say that a cut E separates a set A from a set B if $A \subset E \cup NE$ and $B \subset E^* \cup NE$ or $B \subset E \cup NE$ and $A \subset E^* \cup NE$. Here we also require that neither A nor B is a subset of NE . The separator S is said to separate A from B if for some cut E with $NE = S$, E separates A and B .

Set $a = |E \cap NF|$, $b = |F^* \cap NE|$, $c = |E^* \cap NF|$, $d = |F \cap NE|$ and $m = |NE \cap NF|$, see Figure 5.

Proof for Example 2.4. It follows from the construction that (A1) and (A2) are satisfied. We have show that (A3) is satisfied.

Let E, F be cuts. We want to show that opposite corners have components that are cuts. We know that there are κ -inseparable sets Y_1, Y_2 separated by NE , and κ -inseparable sets Y_3, Y_4 separated by NF . Each Y_i determines a unique corner C_i of E and F , i.e. $Y_i \subset C_i \cup NC_i$. Note that at this stage we cannot rule out the possibility that $Y_i \subset NC_i$ and so has empty intersection with each corner. However there is only one corner for which $Y_i \subset C_i \cup NC_i$, because otherwise Y_i would be contained in NE or in NF .

Two different Y_i may determine the same corner. There must be two Y_i 's which determine opposite corners. For if this is not the case then all four Y_i 's will determine one of a pair of adjacent corners. But then there will either be no Y_i 's separated by NE or no Y_i 's separated by NF . Suppose then, say, that $Y_1 \subset (E \cap F) \cup N(E \cap F)$, $Y_2 \subset (E^* \cap F^*) \cup N(E^* \cap F^*)$.

Consider Figure 5. We see that $|N(E \cap F)| \leq a + m + d$ and $|N(E^* \cap F^*)| \leq b + m + c$. But $\kappa = |NE| = b + m + d = |NF| = a + m + c$, and so $2\kappa = a + b + c + d + 2m$. Hence either $|N(E \cap F)| < \kappa$ or $|N(E^* \cap F^*)| < \kappa$ or $|N(E \cap F)| = |N(E^* \cap F^*)| = \kappa$. But $N(E \cap F)$ and $N(E^* \cap F^*)$ both separate Y_1 and Y_2 and so, by the minimality of κ , $|N(E \cap F)| = |N(E^* \cap F^*)| = \kappa$. It follows that the opposite corners have components that are cuts and so A3 is satisfied. \square

3. MINIMAL SUBSYSTEMS

Let \mathcal{C} be a cut system in a connected graph X . Let κ be the smallest cardinality of a separator. A separator with κ elements is a *minimal separator*. A cut E in \mathcal{C} is called *minimal* if $|NE| = \kappa$.

In Example 2.4 every separator is minimal.




FIGURE 5. Cuts, corners, links and centre

Lemma 3.1. *Let E and F be minimal cuts and assume that $E \cap F$ and $E^* \cap F^*$ are large. Then $N(E \cap F)$ and $N(E^* \cap F^*)$ are minimal separators and $NE \cap NF = N(E \cap F) \cap N(E^* \cap F^*)$.*

In [12, Theorem 2] and [13, Proposition 2.1], Jung and Watkins prove a similar result.

Proof of Lemma 3.1. This is similar to the proof for Example 2.4.

From the diagram, $\kappa = a + m + c = b + m + d$ and hence

$$(1) \quad 2\kappa = a + b + c + d + 2m.$$

Also

$$\begin{aligned} |N(E \cap F)| &= a + d + m_1, \\ |N(E \cap F^*)| &= a + b + m_2, \\ |N(E^* \cap F^*)| &= b + c + m_3, \\ |N(E^* \cap F)| &= c + d + m_4, \end{aligned}$$

where in the above m_1, m_2, m_3, m_4 are the size of subsets of $NE \cap NF$ and so each of these numbers is at most m . For instance,

$$m_1 = |N(E \cap F) \cap NE \cap NF|.$$

The intersections $E \cap F$ and $E^* \cap F^*$ and their $*$ -complements are large. Hence $a + d + m_1 \geq \kappa$ and $b + c + m_3 \geq \kappa$. Now $a + b + c + d + m_1 + m_3 \geq 2\kappa$ and (1) implies $m_1 = m_3 = m$ and $a = b, c = d$. Thus $|N(E \cap F)| = |N(E^* \cap F^*)| = \kappa$. We also get $NE \cap NF = N(E \cap F) \cap N(E^* \cap F^*)$. \square

Corollary 3.2. *The minimal cuts in a cut system form a cut system.*

Proof. It is clear that the minimal cuts and separators in some given cut system \mathcal{C} satisfy Axioms (A1) and (A2), while axiom (A3) follows from Lemma 3.1. \square

A corner of two sets of vertices is called *isolated* if it is small and if its adjacent links are both empty. We call two sets of vertices *nested* if they have an isolated corner.

Lemma 3.3. *The inclusion $E \subset F$ is equivalent to $F^* \subset E^*$ for any sets of vertices E and F . If this inclusion holds for cuts of a cut system, then $E \cap F^*$ is an empty isolated corner.*

If some link of a small corner of two minimal cuts is empty then both links are empty (and the corner is isolated).

Minimal cuts have either (i) no empty link and two large opposite corners or (ii) two non-empty links and two empty links which are adjacent to an isolated corner or (iii) all four links are empty, two opposite corners are large and at least one corner is empty. In case (i) the cuts are not nested and in the cases (ii) and (iii) they are nested.

Proof. Let $E \subset F$. Then $E \cup NE \subset F \cup NF$ and after complementation we get $F^* \subset E^*$. Now $E \cap F^*$ is isolated, because it is empty (the empty set is small), and its links $E \cap NF$ and $F^* \cap NE$ are empty too. Note that $E \cup NE \subset F \cup NF$ does not imply $E \subset F$ in general.

Let E, F be cuts of a cut system. By (A3) two opposite corners are large. We saw in the proof of Lemma 3.1 that the links of a small corner of two minimal cuts must have the same number of elements. Thus if one link is empty then so is the other link. Also if two opposite corners are large then the links of each of the other two corners have the same number of elements. Thus it is not possible for minimal cuts E and F to have exactly one or three empty links. Also two empty links must be adjacent and the corresponding corner must be small. If (iii) all links are empty then one corner has to be empty, otherwise one of the cuts would not be connected. \square

We define a *slice* to be a component of $VX \setminus S$ that is not a cut, where S is a separator. A cut is called an *A-cut* if it is nested with all other minimal cuts. A cut E is called a *B-cut* if E^* has only one large component.

Theorem 3.4. *A minimal cut is either an A-cut or a B-cut.*

Proof. Let E be a minimal cut which is not an A-cut. Then there is a minimal cut F which is not nested with E . We have to show that E^* has only one large component. By case (i) of Lemma 3.3, all links and corners are not empty.

There is an A in $\{E^* \cap F, E^* \cap F^*\}$ which is large and which has a large opposite corner. This corner A has a large component C and Lemma 3.1 implies $|NA| = \kappa$. Now $NC \subset NA$ and $NC \geq \kappa$, because C is a cut. This implies $NC = NA$. Let E_0^* be the component of E^* which contains C . Every vertex x in $E^* \cap NF$ is adjacent to C . Hence x is in E_0^* and $E^* \cap NF \subset E_0^*$.

Suppose there is another large component E_1^* of E^* . Axiom (A1) implies that NE_1^* is a separator. Hence $NE_1^* = NE^*$, otherwise NE_1^* would have less than κ elements. Now $E_1^* \cap E^* \cap NF = E_1^* \cap NF = \emptyset$, because we have seen before that

$E^* \cap NF \subset E_0^*$. There is an element y in $F \cap NE^*$ and an element z in $F^* \cap NE^*$, because all four links are not empty. Since $NE_1^* = NE$ and E_1^* is connected, there is a path from y to z which is completely contained in E_1^* , except for its end-vertices y and z . This path has to intersect $E^* \cap NF$ which contradicts $E_1^* \cap NF = \emptyset$. Hence E^* has exactly one large component. \square

The proof of Theorem 3.4 provides information about the slice components of $VX \setminus NE$ when E is a minimal cut.

Lemma 3.5. *Let \mathcal{C} be a minimal cut system. Distinct slices are disjoint. A slice has empty intersection with each separator. If Q is a slice, then no pair of elements of NQ are separated by any separator.*

Proof. Let Q_1, Q_2 be distinct slice components of $VX \setminus NE$ and $VX \setminus NF$ respectively where E, F are minimal cuts. The argument of Theorem 3.4 shows that if E, F are not nested then NF intersects a single large component of E^* and vice versa and no other component. Thus Q_1 and NF are disjoint. If E, F are nested, then NF has empty intersection with either E or E^* . If it has non-empty intersection with E , then NF has empty intersection with Q_1 . If it has non-empty intersection with E^* , then it follows from (A3) that it has non-empty intersection with some cut component of E^* . By relabeling this component as E we see that Q_1 has empty intersection with NF by the case already considered.

We now consider corners and links of the pair Q_1, Q_2 . We have shown that the links $Q_1 \cap NQ_2, Q_2 \cap NQ_1$ are empty. By the connectedness of Q_1, Q_2 this must mean that either $Q_1 = Q_2$ and all the other corners and links are empty, or $Q_1 \cap Q_2 = \emptyset$. Finally suppose $x, y \in NQ$ for some slice Q and x, y are separated by NE for some cut E . Now Q is connected and the path in Q joining x, y must intersect the separator NE , which we have already shown cannot occur. \square

This lemma enables us to replace a graph X with a cut system of minimal cuts with another graph \hat{X} with essentially the same cut system, but in which there are no slices. In this new cut system two cuts are nested if and only if there is an empty corner with empty adjacent links.

Let, then, X be a graph with a cut system \mathcal{C} of minimal cuts. We define \hat{X} as follows $V\hat{X} \subset VX$ and $v \in V\hat{X}$ if and only if $x \notin Q$ for every slice Q . Note that if $v \in NE$ for any $E \in \mathcal{C}$, then by Lemma 3.5 $v \in V\hat{X}$. Two vertices $u, v \in V\hat{X}$ are joined by an edge in \hat{X} if and only if u, v are joined by an edge in X or $u, v \in NQ$ for some slice Q . We define $\hat{\mathcal{C}}$ to be the set of subsets \hat{E} of $V\hat{X}$, where $\hat{E} = E \cap V\hat{X}$ for some $E \in \mathcal{C}$.

Theorem 3.6. *The set $\hat{\mathcal{C}}$ is a cut system of minimal cuts in \hat{X} . There are no slices in $\hat{\mathcal{C}}$.*

Proof. Let E be a cut in \mathcal{C} . We need to show that $\hat{E} = E \cap V\hat{X}$ is connected in \hat{X} . Any two points of \hat{E} are joined by a path with vertices in E . If this path does not contain any vertices of a slice then it is a path in \hat{E} . If it does contain vertices that are in a slice then delete these vertices. The resulting sequence will be a path in \hat{E} . For if there is a vertex in the original path which is in a slice Q , then there will be vertices u, v in the path, which are also in NE and such that any vertex in the path between u, v will be in Q . These vertices do not lie in any other slice. On deleting these vertices we see that u, v are joined by an edge in \hat{X} . It is easy to deduce all the other properties of a cut system. There will be no slices in \hat{X} . \square

Let \mathcal{C} be a cut system in the graph X . A subset $A \subset VX$ is a *pre-cut* if either A or A^* is a cut.

Let E, F be pre-cuts. We say that E is *almost contained* in F (notation: $E \subset^a F$) if $E \cap F^*$ is an isolated corner. Such a corner is a union of slices. Thus $E \subset^a F$ if and only if $\hat{E} \subset \hat{F}$ in \hat{X} . We analyze the properties of the ordered set (\mathcal{C}, \subset^a) . This will be the same as the properties of $(\hat{\mathcal{C}}, \subset)$ and so we may as well assume that X has no slices and that the order relation is \subset . We know that $E \subset F$ is equivalent to $F^* \subset E^*$. Sets of vertices E and F are nested if and only if one of the conditions $E \subset F$, $E^* \subset F$, $E \subset F^*$ or $E^* \subset F^*$ holds. Also \subset is transitive. It requires a certain amount of effort to prove directly that \subset^a is transitive.

Lemma 3.7. *Let D, E, F be minimal cuts and let E and F be not nested.*

If D is nested with E then it is nested with each cut component of two adjacent corners of E and F . If D is nested with E and F , then it is nested with every cut that is a component of a corner of E and F .

Proof. Suppose D is nested with E . We know that E, F are B -cuts. We can assume that there are no slices by replacing X with \hat{X} if necessary. Thus E^*, F^* are cuts. By changing E to E^* if necessary we can assume that $D \subset E^*$ or $D^* \subset E^*$. If $D \subset E^*$ then $D \subset E^* \cup F$ and $D \subset E^* \cup F^*$ and so D is nested with any cut in the corners $E \cap F^*$ or $E \cap F$. Similarly if $D^* \subset E^*$, then D^* is nested with any cut in the same two corners. But D is nested with a cut if and only if D^* is nested with the same cut. A similar argument gives the second statement. \square

4. SEPARATION BY FINITELY MANY CUTS

We recall some results concerning separating by removing sets of edges. A *minimal k -edge separator* is a set of k edges whose complement is disconnected and the complement of any proper subset of this separator is connected. Note that the complement of a minimal edge separator has exactly two components and each edge in the separator is adjacent to both components. Every finite set of edges with a disconnected complement contains a minimal edge separator. Let p be any edge of some connected graph and let k be any integer. Thomassen and Woess have given a short proof of the fact that there are only finitely many minimal k -edge separators containing p , see [24, Proposition 4.1].

We call a set S of k vertices a *minimal k -separator*, if $VX \setminus S$ has at least two components which are adjacent to all elements of S . That is, S is minimal with respect to separating these two components. Note that this is a general graph-theoretic definition which does not refer to our axiomatic cut systems and their separators. Also note that a graph may have minimal k -separators for different values of k . We say that a minimal k -separator S *separates* vertices u and v *properly* if u and v lie in distinct components of $VX \setminus S$ which are adjacent to every element of S . When we think of our cut systems satisfying the axioms (A1)-(A3) then we observe that minimal separators S are κ -vertex separators which separate vertices in distinct cuts properly whenever the boundaries of these cuts are S .

Thomassen and Woess showed that in a locally finite graph, there are only finitely many minimal k -vertex separators containing a given vertex x , see [24, Proposition 4.2]. This does not hold in non-locally finite graphs. But we can prove

the following two lemmas for finite vertex separators in non-locally finite graphs using similar arguments as in the proof of [24, Proposition 4.1].

Lemma 4.1. *Every pair of vertices in a connected graph is separated properly by only finitely many k -vertex separators, for any given k .*

Proof. We use induction on k . Any minimal 1-separator of vertices x and y is a vertex in any path joining x and y and so the lemma is true for $k = 1$.

Suppose the lemma holds for all minimal k -separators in all connected graphs. Let x and y be any vertices of a graph X . We choose a path π from x to y and assume that there are infinitely many minimal $(k + 1)$ -separators, $k + 1 \geq 2$, separating x from y properly. Then there is a vertex $z \in \pi \setminus \{x, y\}$ which is contained in infinitely many of these $(k + 1)$ -separators. If S_1 and S_2 are such $(k + 1)$ -separators then $S_1 \setminus \{z\}$ and $S_2 \setminus \{z\}$ are distinct k -separators of x and y in $X - \{z\}$, and they separate x from y properly in $X - \{z\}$. Hence there are infinitely many distinct k -separators in $X - \{z\}$ which separate x from y properly, in contradiction to the induction hypothesis. \square

Lemma 4.2. *A minimal pre-cut is nested with all but finitely many minimal pre-cuts.*

Proof. Let E and F be minimal pre-cuts which are not nested. By Lemma 3.3 all links are not empty. Hence NF has to separate two elements x and y of NE , and NF separates these vertices properly as a κ -vertex separator. Lemma 4.1 says that there are only finitely many such separators NF . Since NE is finite and since, by Theorem 3.4, the complement $VX \setminus NF$ of each such set NF has exactly two large components (i.e. NF corresponds to two cuts), we obtain the statement of the lemma. \square

5. OPTIMALLY NESTED CUTS

Let \mathcal{C} be a cut system of minimal cuts. Let C be a cut and let $M(C)$ be the set of cuts which are not nested with C . Set $m(C) = |M(C)|$. It follows from Lemma 4.2 that $m(C)$ is finite. If $E \in \mathcal{C}$, then each large component of E^* is a cut. Put $m(E^*) = m(C)$ where C is a large component of E^* . There is no ambiguity in doing this as if there are two cuts C_1, C_2 which are large components of E^* then E, C_1, C_2 are all A -cuts and $m(C_1) = m(C_2) = 0$.

Lemma 5.1. *Let E and F be B -cuts which are not nested, and suppose $E \cap F$ and $E^* \cap F^*$ are large, then*

$$m(E \cap F) + m(E^* \cap F^*) < m(E) + m(F).$$

Proof. Let D be any minimal cut. If D is in $M(E^* \cap F^*) \cap M(E \cap F)$ then, by Lemma 3.7, D is in $M(E)$ and in $M(F)$. Hence if D is counted twice on the left of the above inequality then it is also counted twice on the right.

If otherwise D is in $M(E \cap F) \setminus M(E^* \cap F^*)$ or in $M(E^* \cap F^*) \setminus M(E \cap F)$, that is D is counted exactly once on the left, then, again by Lemma 3.7, D is in $M(E)$ or in $M(F)$. Hence D is counted at least once on the right side of the inequality. We have now proved that $m(E \cap F) + m(E^* \cap F^*) \leq m(E) + m(F)$. Since $E \in M(F)$ and $F \in M(E)$, the cuts E and F are counted on the right side, but not on the left side, and we see that this inequality is a strict inequality. \square

Set $m = \min\{m(E) \mid E \text{ is a minimal cut.}\}$. This minimum exists, because the values $m(E)$ are all finite. A minimal cut E with $m(E) = m$ is called *optimally nested*. Every non-empty cut system contains an optimally nested minimal cut.

Theorem 5.2. *Every optimally nested minimal cut is nested with all other optimally nested minimal cuts. The optimally nested minimal cuts form a nested cut system.*

Proof. Suppose there are optimally nested minimal cuts E and F which are not nested. By (A3) there will be opposite corners that are large. By relabeling we can assume that these large corners are $E \cap F$ and $E^* \cap F^*$ and by Lemma 3.1, each of $E \cap F$ and $E^* \cap F^*$ has a component which is a minimal cut. Now Lemma 5.1 says that

$$m(E \cap F) + m(E^* \cap F^*) < m(E) + m(F) = 2m.$$

Thus one of the summands on the left side is less than m , contradicting the minimality of m . \square

6. ORDERS AND TREES

We will show that the cut system of minimal nested cuts can be regarded as the edge set of directed tree. First we consider again the structure tree given by cut-points and 2-blocks described in the introduction. In the discussion here we do allow disconnecting edges. The vertices of a disconnecting edge form a 2-block.

Example 6.1. Let X be a connected graph. If X is not 2-connected then it has *cut-points*, i.e. vertices whose removal disconnects the graph. If this happens, then VX decomposes into a collection of maximal 2-inseparable subsets or *2-blocks*. Any two 2-blocks intersect in at most one vertex and this vertex will be a cut-point. Every edge of X joins vertices in exactly one 2-block. Note that the boundary of a 2-block in X consists of cut-points. Let \mathcal{C} be the set of connected sets of vertices E such that $NE = \{x\}$ and x is a cut-point. Then \mathcal{C} is a system of minimal nested cuts.

As noted in the introduction, associated with this decomposition is a tree T in which $VT = \mathcal{B} \cup \mathcal{S}$, where \mathcal{S} is the set of cut-points, \mathcal{B} is the set of 2-blocks and there is an edge joining $b \in \mathcal{B}$ with $s \in \mathcal{S}$ if and only if $s \in b$. But note that the edges of this tree can be regarded as the set \mathcal{C} . Thus for $E \in \mathcal{C}$, let $t(E) = NE$ and let $o(E)$ be the block that is contained in E that has the cut-point NE in its boundary. If we direct the edges of T so that an arrow points from $o(E)$ to $t(E)$, then the vertices in \mathcal{S} have every adjacent edge pointing towards if and every vertex in \mathcal{B} has every adjacent arrow pointing away from it. Any path in T will have alternating arrows as one proceeds along it.

In general let \mathcal{C} be a nested cut system of minimal cuts. By replacing X by \hat{X} if necessary, we can assume that there are no slices and so \subset^a is the same as \subset for pre-cuts E, F . Let \mathcal{S} be the set of separators. We define \mathcal{B} as the set of equivalence classes for a particular equivalence relation \sim on \mathcal{C} . Thus for $E, F \in \mathcal{C}$ put $E \sim F$ if either

- (i) $E = F$ or
- (ii) $NE \neq NF$ and $E^* \subset F$ and if $E^* \subset D \subset F$, for $D \in \mathcal{C}$, then $D = F$, and if $E^* \subset D^* \subset F$, for $D \in \mathcal{C}$, then $E = D$.

We prove that \sim is an equivalence relation. Clearly \sim is reflexive. It is symmetric, because $E^* \subset F$ if and only if $F^* \subset E$, and $E \subset F$ if and only if $E^* \subset F^*$ for $E, F \in \mathcal{C}$.

Suppose $D \sim E$ and $E \sim F$ and suppose that D, E, F are distinct elements of \mathcal{C} , so that $D^* \subset E$ and $E^* \subset F$. We know that D, F are nested, so that there is one of four possibilities. If (1) $D \subset F^*$ then $E^* \subset D$ implies $E^* \subset F^*$. Since $F^* \subset E$ we get $E^* \subset E$, which is impossible. If (2) $D \subset F$ then, $F^* \subset D^* \subset E$ which implies $F^* = D^*$, and hence $F = D$, which we have excluded. If (3) $D^* \subset F^*$ then $D^* \subset F^* \subset E$ and either so $D = F$ or $F^* = E$. We have excluded $D = F$ and if $F^* = E$ then $NF = NE$ and so $E = F$ which is also excluded. If (4) $D^* \subset F$ then either $D \sim F$ or there is an $A \in \mathcal{C} \cup \mathcal{C}^*$, $A \neq F, A \neq D^*$ such that

$$D^* \subset A \subset F.$$

We have again four cases, since A and E are nested. If (1) $E \subset A$ then $E \subset F$. Since $E \sim F$, we have $E^* \subset F$. Since $F^* \subset E^*$ we now get $F^* \subset F$ which is impossible. If (2) $E^* \subset A$ then $E^* \subset A \subset F$ and $E^* = A$. Then $D^* \subset E^*$. Together with $E^* \subset D$, this would imply $D^* \subset D$ which again is impossible. If (3) $E \subset A^*$ then $A \subset E^* \subset D$, which implies $D^* \subset A^*$, in contradiction to $D^* \subset A$. If (4) $E^* \subset A^*$ then $D^* \subset A \subset E$ and so $A = E$, by the definition of \sim . This implies $E \subset F$, contradicting $F \subset E^*$.

We obtain a directed graph $T = T(\mathcal{C})$

$$VT = \mathcal{S} \cup \mathcal{B} \quad \text{and} \quad ET = \mathcal{C}.$$

where \mathcal{S} is the set of separators and $\mathcal{B} = \mathcal{C}/\sim$, where $t(E) = NE$ of $E \in \mathcal{C}$ and $o(E)$ is the \sim -class which contains E . Clearly T is a bipartite graph since $o(E)$ and $t(E)$ are in disjoint sets. In particular there are no loops. We will shortly show that T is a tree. If we direct the edges of T so that an arrow points from $o(E)$ to $t(E)$, then each vertex in \mathcal{S} has every adjacent edge pointing towards it and each vertex in \mathcal{B} has every adjacent arrow pointing away from it. Any path in T will have alternating arrows as one proceeds along it.

Lemma 6.2. *The graph $T = T(\mathcal{C})$ is a tree.*

Proof. Suppose the underlying graph of T contains a cycle. The graph T contains an alternating closed cycle of length $2n$ for some $n \geq 2$. There are elements E_1, E_2, \dots, E_{2n} of $\mathcal{C} \cup \mathcal{C}^*$ corresponding to the edges of this directed cycle such that

$$E_1 \subset E_2 \subset E_3 \dots E_{2n} \subset E_1 \subset E_2 \subset \dots,$$

in contradiction to \subset being antisymmetric.

Let E and F be elements of \mathcal{C} . If $E \subset F$, then Lemma 4.1 implies that there is a finite path in T connecting $o(E)$ and $t(F)$. Hence T is connected. \square

The following example, illustrated in Figure 6, shows that it is necessary to use \subset^a rather than \subset when defining nestedness if there are slices.

Example 6.3. Set $VX = \mathbb{Z} \cup \{o, i\}$ and

$$EX = \{\{x, x+1\} \mid x \in \mathbb{Z}\} \cup \{\{x, o\} \mid x \in \mathbb{Z}\} \cup \{\{o, i\}\}.$$

The minimal number of vertices needed to separate the two ends of X is 2. The (connected) minimal cuts are of the form $E_k^+ = \{k+1, k+2, \dots\}$ or $E_k^- = \{k-1, k-2, \dots\}$, where $NE_k^+ = NE_k^- = \{k, o\}$. The group of automorphisms acts transitively on the cut system $\mathcal{C} = \{E_k^+, E_k^- \mid k \in \mathbb{Z}\}$. Also \mathcal{C} is the only minimal

end separating automorphism invariant cut system. If $l > r$ then E_l^- and E_r^+ do not have an empty corner but they are nested because $\{i\}$ is their isolated corner. The set $\{i\}$ is the only slice. The graph \hat{X} is obtained by deleting i .

The tree $T(\mathcal{C})$ is a double ray.




FIGURE 6. Structure tree for a two-ended graph

It is possible to define the set \mathcal{B} in a different - possibly better - way. Let \mathcal{C} be a minimal cut system. For the moment we do not assume \mathcal{C} is nested. A subset Y of VX is said to be \mathcal{C} -inseparable if for every $A \in \mathcal{C}$ either $Y \subset A \cup NA$ or $Y \subset A^* \cup NA$ but not both. It follows from Zorn's Lemma that every \mathcal{C} -inseparable set is contained in a maximal \mathcal{C} -inseparable set. Thus if one has an increasing sequence of \mathcal{C} -inseparable sets $B_1 \subset B_2 \subset \dots$ and $B = \bigcup_n B_n$, then B is \mathcal{C} -inseparable, since if $A \in \mathcal{C}$ and k is a positive integer such that NB_n has more than NA elements if $n \geq k$, then for all $n \geq k$ either $B_n \subset A \cup NA$ or $B_n \subset A^* \cup NA$ and the same is true for B . A maximal \mathcal{C} -inseparable set B is called a \mathcal{C} -block, if B is not contained in any slice. Thus a \mathcal{C} -block corresponds to a $\hat{\mathcal{C}}$ -block in \hat{X} .

For a nested minimal cut system, we define \mathcal{B} to be the set of \mathcal{C} -blocks.

Lemma 6.4. *If $A \in \mathcal{C}$ then there is a unique $B \in \mathcal{B}$ such that $NA \subset B$ and $B \subset NA \cup A$.*

Proof. By changing to \hat{X} if necessary we can assume that there are no slices.

Let S be a separator. Thus $S = NA$ for at least two components A of $X \setminus S$ and they are cuts. If S is the only separator, then each $B \in \mathcal{B}$ is a set $B = A \cup NA$ where $A \in \mathcal{C}$ is such a component. In general let \sim be the equivalence relation on \mathcal{C} defined above. Let

$$I = \bigcap_{E \sim A} (E \cup NE).$$

It is fairly easy to check that I is \mathcal{C} -inseparable and satisfies the required conditions \square

Put $i(A) = B$, where B is as in Lemma 6.4. In fact every block $B = i(A)$ for some $A \in \mathcal{C}$. To see this note that $B \subset A \cup NA$ for some A . Now choose A so that it is minimal with this property.

7. STRUCTURE TREES

Let G be a group acting on a connected graph X and let \mathcal{C} be a nested cut system of minimal cuts, invariant under G . Let $T = T(\mathcal{C})$ be the tree constructed in the last section. This will be a G -tree, i.e. the action of G on \mathcal{C} induces an action of G on T .

We now show how a ray in X determines either a unique end or a unique block vertex of T . Let R be a ray in X . For each cut E , R is eventually in E or E^* . Thus for each pair E, E^* , the ray determines a choice of just one of E or E^* . If R eventually lies in D and $D \subset E$ then clearly R eventually lies in E . It follows that R determines an orientation of the edges of the tree T . This orientation is different from the orientation given previously in which paths have alternating arrows. In this new orientation $d \leq e, e \leq f$ implies $d \leq f$. Under this orientation all the edges point to a single vertex or end of T . If they point to a vertex, then this vertex cannot be a separator, since the vertices of a ray are distinct, and so it can only visit a separator S at most $|S|$ times. The ray determines the block $B \in \mathcal{B}$ if and only if it contains infinitely many vertices of B . We say that the ray R *belongs to* the block vertex or end of T .

Now we show that if two rays belong to the same end of T then they belong to the same end of X .

Let R_1, R_2 belong to the same end of T . Thus there is an infinite sequence of cuts $E_1 \supset E_2 \supset E_3 \dots$ such that both R_1, R_2 are eventually in any one of the E_i 's. Note that the intersection of the E_i 's is empty. For if $u \in VX$ is in every E_i , then it will be separated from a vertex in E_1^* by every NE_i . This contradicts Lemma 4.1. Suppose there is a cut D which separates R_1, R_2 , say R_1 is eventually in D and R_2 is eventually in D^* . For each i , both $E_i \cap D$ and $E_i \cap D^*$ contain rays. For only finitely many i 's NE_i separates vertices in ND . In fact either $NE_i \subset D$ for i large enough or $NE_i \subset D^*$ for i large enough. But if, say, $NE_i \subset D$ for i large enough, then both R_1, R_2 are eventually in D , since both rays contain vertices from infinitely many of the NE_i 's.

Let \mathcal{C} be a cut system. If a set Y is \mathcal{C} -inseparable, then it is \mathcal{C}' -inseparable for any subsystem \mathcal{C}' of \mathcal{C} . Thus any \mathcal{C} -block is contained in a \mathcal{C}' -block. The inclusion may be proper, and there may be \mathcal{C}' -blocks which contain no \mathcal{C} -block. Thus in Figure 2 the set B consisting of the four vertices of the central dotted 4-cycle do not form a \mathcal{C} -block for the cut system \mathcal{C} of all cuts A with $|NA| = 2$. But if we restrict to the nested cut system \mathcal{C}' of the cuts for which NA consists of two adjacent vertices of the 4-cycle, then B is a \mathcal{C}' -block.

Let \mathcal{C} be a cut system of minimal cuts.

Lemma 7.1. *There is a nested subsystem \mathcal{E} of minimal cuts invariant under G such that distinct \mathcal{C} -blocks lie in distinct \mathcal{E} -block and any two rays that are separated by an element of \mathcal{C} are separated by a cut in \mathcal{E} and any \mathcal{C} block that is separated from an end of X by a minimal cut can be separated from that end by a cut in \mathcal{E} .*

Proof. We know that the set of optimally nested cuts is a nested sub-system \mathcal{N} of cuts in \mathcal{C} invariant under G . If this cut system does not have the required properties listed in the statement of the lemma, then we show that we can enlarge it, and obtain a bigger sub-system \mathcal{N}' with the same properties. In fact all we will assume about \mathcal{N} is that it is a nested subsystem of \mathcal{C} invariant under G and that there is an integer $M \geq m$ such that $m \leq m(A) \leq M$ for every $A \in \mathcal{N}$ and no cut

A in \mathcal{C} with $m(A) < M$ separates two \mathcal{C} blocks, or separates two ends or separates a \mathcal{C} -block and an end unless they are already separated by a cut in \mathcal{N} . Suppose then that we have such an \mathcal{N} but that it does not have the required properties listed for \mathcal{E} in the statement of the theorem.

Then one of the following holds.

- (a) There are distinct \mathcal{C} -blocks B_1, B_2 that lie in the same \mathcal{N} -block.
- (b) There are two rays in X that are separated by some $A \in \mathcal{C}$ that are not separated by any $E \in \mathcal{N}$.
- (c) There is a ray and a \mathcal{C} -block B_1 that are separated by some $A \in \mathcal{C}$ that are not separated by any $E \in \mathcal{N}$.

Suppose that (a) holds and that the minimal value M for which there is a cut $A \in \mathcal{C}$ with $m(A) = M$ is less than or equal to the corresponding value for a cut A satisfying (b) or (c). Let B be the \mathcal{N} -block that contains B_1, B_2 . For some $A \in \mathcal{C}$, $B_1 \subset A \cup NA$ nor $B_2 \subset A^* \cup NA$, i.e. NA separates B_1 and B_2 . Choose A with this property so that $m(A)$ is minimal.

We show first that A is nested with every $E \in \mathcal{N}$. For if it is not nested with E , we know that $B \subset E \cup NE$ or $B \subset E^* \cup NE$. Without loss of generality assume $B \subset E \cup NE$. Now either $E \cap A$ or $E \cap A^*$ is in \mathcal{C} and its boundary separates B_1 and B_2 by Lemma 5.1 we have $m(A \cap E) + m(A^* \cap E^*) < m(A) + m(E) \leq m(A) + M$. Since $m(A^* \cap E^*) \geq M$ we have $m(A \cap E) < m(A)$ which is a contradiction. Thus A is nested with every $E \in \mathcal{N}$.

It remains to show that A is nested with gA for every $g \in G$. We can then enlarge \mathcal{E} by adding all the translates $gA, g \in G$. Suppose A and gA are not nested. It is not hard to see that there are components of opposite corners each of which will separate two of B_1, B_2, gB_1, gB_2 . Let these corners be $A \cap gA, A^* \cap gA^*$. By Lemma 5.1 we have $m(A \cap gA) + m(A^* \cap gA^*) < m(A) + m(gA) = 2m(A)$. By the choice of A we have a contradiction.

A similar argument works if the minimum value for an $m(A)$ which separates two \mathcal{C} blocks not separated by a cut in \mathcal{N} is greater than that of an A satisfying (b) or (c). We just replace the cuts separating blocks by cuts separating ends, or by cuts separating an end and a block. □

We have proved the following:

Theorem 7.2. *Let X be a connected graph with automorphism group G and cut system \mathcal{C} . There is a nested sub-system of minimal cuts \mathcal{E} in $\mathcal{C}(X)$ invariant under G with the following properties. If two rays are separated by a minimal cut $D \in \mathcal{C}$ then they are separated by a cut $E \in \mathcal{E}$. If two \mathcal{C} -blocks are separated by a minimal cut $E \in \mathcal{C}$ then they are separated by a cut in \mathcal{E} . If an end of X and a \mathcal{C} -block are separated by a minimal cut in \mathcal{C} , then they are separated by a cut in \mathcal{E} .*

Let $T = T(\mathcal{E})$ be the associated G -tree. A ray (which belongs to an end of X) will either belong to an end or a vertex of T . Two rays which determine the same end of X will belong to the same end or the same vertex of T . Two rays which belong to different ends of X will belong to the same end or the same vertex of T if and only if they cannot be separated by a minimal cut D . A \mathcal{C} -block belongs to a unique vertex of T . Distinct blocks lie in the same vertex if and only if they cannot be separated by a minimal cut in \mathcal{C} .

As in [2] one can obtain a stronger result in which the nested sub-system does not only contain minimal cuts. We describe how to do this. Suppose there are two \mathcal{C} -blocks that lie in the same \mathcal{E} -block. Two such blocks will be separated by some NA for $A \in \mathcal{C}$. In fact they will be separated by an NA which is nested with every $E \in \mathcal{E}$. This can be proved by choosing such an A for which the number of E with which it is not nested is minimal, and then using Lemma 5.1, we can show that this number is zero. For such an A , NA lies in a unique \mathcal{E} -block B .

The set of cuts

$$\mathcal{C}_B = \{A \in \mathcal{C} \setminus \mathcal{E} \mid A \text{ nested with every } E \in \mathcal{E}, NA \subset B\}$$

is a cut system. It is easy to see that (A1) and (A2) are satisfied.

For (A3), let $C, D \in \mathcal{C}_B$. Since $E, F \in \mathcal{C}$ they have opposite corners that have components in \mathcal{C} . We want to show that for each such corner at least one of the corners is in \mathcal{C}_B . For each $E \in \mathcal{E}$, either $B \subset E \cup NE$ or $B \subset E^* \cup NE$. Consider the set $\mathcal{E}_B = \{E \in \mathcal{E} \mid NE \subset B, B \cap E = \emptyset\}$. In the structure tree $T(\mathcal{E})$, the set \mathcal{E}_B consists of the edges with $t(E) = B$. We know that each $C \in \mathcal{C}_B$ is nested with each $E \in \mathcal{E}_B$. In fact we have either $E \subset C$ or $E \subset C^*$, since both $E^* \cap C$ and $E^* \cap C^*$ have non-trivial intersections with \mathcal{C} -blocks. Since the same is also true for D it follows that for every $E \in \mathcal{E}_B$ either $E \subset C \cap D$ or $E \subset C^* \cup D^*$ and so E is nested with each component of each corner of C and D . It follows fairly easily that if $E \in \mathcal{E}$ (not just \mathcal{E}_B) then E is nested with each component of a corner of C and D . We need also to show that opposite corners of C, D have components that are not in \mathcal{E} but which are in \mathcal{C} . Since C is connected, neither $C \cap D$ or $C \cap D^*$ can have every component that is in \mathcal{E}_B . For if this were the case then the link between the two corners would be empty. By replacing C by a component of C^* we can in fact conclude that no corner of C, D can have every component in \mathcal{E}_B . But by (A3) for \mathcal{C} there are opposite corners that have components in \mathcal{C} . Thus (A3) is verified for \mathcal{C}_B .

We can now enlarge \mathcal{E} by adding a nested subsystem of minimal cuts in \mathcal{C}_B . If $g \in G, gB \neq B$, then one can show fairly easily that any cut in \mathcal{C}_B is nested with any cut in \mathcal{C}_{gB} and so we can further enlarge \mathcal{E} and obtain a new nested cut system which is a G -set.

Using a similar argument we can also enlarge \mathcal{E} if there are two ends of X that are separated by a cut in \mathcal{C} but not by a cut in \mathcal{E} or if there is a \mathcal{C} -block and an end of X that are separated by a cut in \mathcal{C} but not by a cut in \mathcal{E} . Note that any slice in the \mathcal{C}_B will already have been a slice in \mathcal{C} .

Note also that the nested sub-system \mathcal{E} of \mathcal{C} is unique. This is because the construction above is canonical. Thus we start by taking all the optimally nested minimal cuts in \mathcal{C} and then extend by taking all the cuts which are nested with the previous subsystem and which satisfy some minimality condition. This means that the trees T_n are unique, i.e. they depend only on the cut system \mathcal{C} .

We obtain the following strengthened version of Theorem 7.2.

Theorem 7.3. *Let X be a connected graph with automorphism group G and cut system \mathcal{C} . There is a unique nested sub-system \mathcal{E} of cuts in $\mathcal{C}(X)$ invariant under G with the following properties. If two rays are separated by a cut $D \in \mathcal{C}$, then they are separated by a cut $E \in \mathcal{E}$ with $|NE| \leq |ND|$. If two \mathcal{C} -blocks are separated by a cut $D \in \mathcal{C}$, then they are separated by a cut in \mathcal{E} with $|NE| \leq |ND|$. If an end of*

X and a \mathcal{C} -block are separated by a cut in \mathcal{C} , then they are separated by a cut in \mathcal{E} with $|NE| \leq |ND|$.

Let \mathcal{E}_n be the sub-system of \mathcal{E} consisting of cuts $E \in \mathcal{E}$ with $|NE| \leq n$. Then there is a G -tree $T_n = T_{\mathcal{E}_n}$ associated with this cut system. A ray (which belongs to an end of X) will either belong to an end or a vertex of T_n . Two rays which determine the same end of X will belong to the same end or the same vertex of T_n . Two rays which belong to different ends of X will belong to the same end or the same vertex of T_n if and only if they cannot be separated by a cut $D \in \mathcal{C}$ with $|ND| \leq n$. A \mathcal{C} -block belongs to a unique vertex of T_n . Distinct blocks lie in the same vertex if and only if they cannot be separated by a cut $D \in \mathcal{C}$ with $|ND| \leq n$.

We thank Matthias Hamann for pointing out an error in the proof of this theorem in an earlier version of this paper.

It is necessary to restrict to cuts $E \in \mathcal{E}$ with $|NE| \leq n$ in order to get that there are only finitely many edges between two vertices of the associated tree. The whole set \mathcal{E} can be regarded as the directed edge set of a generalized tree called a protree (see [5]).

Each \mathcal{C} -block will be a \mathcal{E} -block. There may be \mathcal{E} -blocks that are not \mathcal{C} -blocks, as can be seen from the example in the introduction. If we join any two vertices by an ideal edge if they are not separated by any cut in the final cut system \mathcal{E} then it can be seen that \hat{X} embeds in a graph in which there is a nested cut system, which is essentially \mathcal{E} and every \mathcal{E} -block is the vertex set of a complete subgraph. This is illustrated in Figure 3.

This last theorem only gives us extra information when we have a cut system that contains cuts which are not minimal. For example it applies to the cut systems Example 2.1 and Example 2.2. It does not give extra information for the cut system Example 2.4. However, as explained in the next section, we can obtain extra information about the connectivity of finite graphs by working with the graphs obtained from the blocks of this latter cut system.

8. APPLICATIONS

It is a consequence of our main results (Theorem 7.2 and Theorem 7.3) that if a G -graph has a non-trivial cut system \mathcal{C} in which the size of separators is bounded, then there is a homomorphism of G to the automorphism group of a tree. Thus there is a G -tree $T = T(\mathcal{E})$ associated with a nested sub-system \mathcal{E} of \mathcal{C} . The actions of groups on trees are completely described in the theory of Bass and Serre (see [20], [19] or [2]). The action of a group G on a tree T is said to be *trivial* if G fixes a vertex. If T has finite diameter - in particular if T is finite - then the action is trivial. The quotient graph $G \backslash T = Q$ is a tree.

For a finite graph X , there is a natural cut system \mathcal{C} . Namely the cut system of Example 2.4. This will then contain a unique nested sub-system \mathcal{E} which will be the directed edge set of a structure tree $T(\mathcal{E})$. We summarize this in the following theorem.

Theorem 8.1. *Let X be a finite graph with automorphism group G . Let κ be the smallest value of k for which there are a pair of k -inseparable subsets of VX which can be separated by a k -separator. There is a unique nested cut system \mathcal{E} of κ -separators, invariant under G with the following property. If two κ -inseparable*

sets are separated by a set of κ vertices in VX , then they are separated by a cut in \mathcal{E} .

The G -tree associated with this nested cut system is called the *structure tree* for X . There may be slices for this cut system. (But we have not been able to construct an example where there are.) If we remove vertices in any slice and join each pair of vertices by an ideal edge if the vertices lie in the same block, then we obtain a G -graph with the same structure tree in which the blocks are complete subgraphs.

If there are k -inseparable sets that are separated by a set of k -elements but are not separated by a separator in \mathcal{E} , then these sets must belong to the same \mathcal{E} -block. Let κ' be the smallest value of k for which this is true. Then $\kappa' > \kappa$. Thus we have κ' -inseparable sets which lie in the same \mathcal{E} -block B and which are separated by a set with κ' elements of VX . By adding (ideal) edges joining any pair of vertices in the same \mathcal{E} -separator we can make the make B into the vertices of a connected graph X_B . Also κ' inseparable sets which are separated by κ' vertices in X will also be separated by κ' vertices in X_B , and so we can repeat our theory and obtain a structure tree T_B for X_B .

We can expand the vertex v_B corresponding to B of the structure tree for \mathcal{E} and get a new expanded structure tree. The expansion involves choosing, for each edge e of $T(\mathcal{E})$ incident with v_B , a vertex of the tree T_B to which e will be attached in the expanded tree. There is not usually a canonical way of doing this. In the case of a finite graph it can be done so that the expanded tree is G -invariant. In the situation of an infinite graph such as in Figure 6 it may not be possible to carry out the expansion and get a G -tree.




FIGURE 7. Structure trees for a finite graph

What happens for a finite graph is illustrated in Figure 7. The graph X shown has automorphism group G of order two, and has one cut point. This gives rise to a cut system \mathcal{E} with two cuts one of which is the top vertex and the other the bottom 5 vertices. These cuts are nested and the corresponding tree ($T(\mathcal{E})$) has three vertices, the middle vertex corresponds to the separating cut point and the two edges correspond to the cuts. There are two blocks, corresponding to the two other vertices of T . The vertices of one block are the vertices of the separating edge of X and the other block B corresponds to X with the top vertex removed. Now B has a nested system of cuts in which there three separators each with two elements. The tree T_B corresponding to this cut system is shown at the bottom right of

the diagram. It has 7 vertices, three of which correspond to these separators and the remaining four vertices correspond to the four 3-cycles in X . We can expand the vertex corresponding to B in $T(\mathcal{E})$ so that we get a structure tree for X , by attaching the dashed edge as shown. This is the only way to attach the edge to preserve the symmetry of the tree. Note though that it involves joining vertices corresponding to separators in the different cut systems, and so one cannot always expand blocks so that a 2-colouring of the expanded tree gives a natural way of getting separators and blocks. The reason one can always expand the tree in the case of a finite graph and obtain a tree on which the automorphism group G acts is because a group action on a finite tree is trivial. There is always a vertex that is fixed by G . Thus one can always attach edges to the vertex in the structure tree for the block that is fixed by the automorphism group of the block, and this will mean that the expanded tree admits the automorphism group G of X .

In this example there is only one vertex in T_B that is fixed by the automorphism group (of order two) but this will not always be the case. Thus the automorphism group may be trivial, in which case any vertex of T_B can be chosen for the expansion, which will not then be canonical.

If a structure tree has infinite diameter, then G may induce a non-trivial action on the structure tree.

We briefly describe the relevant results from Bass-Serre theory.

A group is said to *split* over a subgroup H if it is either a free product of two groups with amalgamation over H , where these two groups contain H as a subgroup of index at least two, or if G is an HNN-extension of some group over H .

Suppose a group G acts transitively and without inversion on the set of edges of a tree T . Then either the quotient T/G is a loop and G is an HNN-extension of the stabilizer of some vertex of T over the stabilizer of an incident edge. Or G has two orbits on VT , the quotient T/G is a graph with two vertices connected by an edge (called a *segment*), and G is a free product of the stabilizers of two adjacent vertices in T with amalgamation over the stabilizer of the edge which connects them. This decomposition is trivial if and only if the stabilizer of an edge is the same as the stabilizer of one of its vertices and the whole group stabilizes the other vertex v . If this happens then the tree T has diameter two, with central vertex v . The action is non-trivial if and only if for each edge $e \in ET$ both components of $T \setminus \{e\}$ contain at least one edge (or equivalently, at least two vertices). In fact if the action is non-trivial, then both components of $T \setminus \{e\}$ are infinite. Thus if G acts transitively without inversion on the set of edges of T then either T has diameter two or G splits over the stabilizer of an edge. The latter happens if and only if for some edge e both components of $T \setminus \{e\}$ intersect the orbit of e .

More generally the action (without inversion) of a group G on a G -tree T is non-trivial if and only if either G splits over an edge stabilizer or it is a strictly ascending union

$$G = \bigcup_n G_n,$$

where $G_1 \subset G_2 \subset \dots$ is an infinite sequence of proper subgroups of G each of which stabilizes an edge of T .

If G is a group, a Cayley graph for G is a connected G -graph with one orbit of vertices and on which G acts freely. The edge orbits will correspond to a set of generators for G . There is a locally finite Cayley graph if and only if G is finitely

generated. Different locally finite Cayley graphs of a finitely generated group are quasi-isometric. The number of ends of a locally finite graph is a quasi-isometry invariant and hence it does not depend on the finite set of generators. Thus we define the *number of ends of a finitely generated group* as the number of ends of its locally finite connected Cayley graphs.

The following was proved by Stallings in a series of papers, see [21, 22, 23].

Theorem 8.2 (Stallings’ Structure Theorem). *A finitely generated group has more than one end if and only if it splits over a finite subgroup.*

The first author generalized Stallings’ theorem in [4] by proving that the cut system of edge cuts, see Example 2.2, has a nested subsystem. We have proved that any cut system has a nested subsystem which separates ends. Thus we have a new proof of the main result in [4]. What we obtain gives a new and relatively simple proof of Stallings’ Structure Theorem, see [16].

There are different ways of generalizing Stallings theorem. One option is to drop the assumption of G being finitely generated. The other options is to consider splittings of finitely generated groups over groups which are not necessarily finite.

There are several ways of how to define ends for non-locally finite graphs, see [14]. The same holds for infinitely generated groups, where we have the further difficulty that without additional assumptions the Cayley graphs of an arbitrary group are not necessarily quasi-isometric. But whenever one defines ends of non-locally finite graphs then in locally finite graphs this definition should yield Freudenthal’s end compactification for locally compact space (see [9, 10, 11]).

One way goes back to Freudenthal and D.E.Cohen [1] and says that G has more than one end if there is a subset A for which A and the complement $G \setminus A$ are both infinite and the symmetric difference of A and Ag is finite for all g in G .

It follows from the Almost Stability Theorem [2] that a group G has more than one end in this sense if and only if G splits over a finite subgroup or G is countably infinite and locally finite. This is a generalization of Stallings’ structure theorem, because in the finitely generated case the definition above is equivalent to all other definitions of ends of graphs and groups. A more revealing way of stating this result follows from the Bass-Serre theory discussed above. Thus a group has more than one end if and only if it has a non-trivial action on a tree with finite edge stabilizers.

For a group that is not finitely generated there is no obvious way to choose a generating set to construct a Cayley graph. If we take the whole group as a set of generators, then the Cayley graph is essentially a complete graph which will have one end in any definition.

Stallings’ theorem can be formulated as “A finitely generated group has a Cayley graph with more than one end if and only if it splits over a finite subgroup.” Here we can just drop the assumption that the group is finitely generated.

Theorem 8.3. *A group has a Cayley graph with more than one end if and only if it splits over a finite subgroup.*

Proof. Suppose G splits over a finite group H . There are two possibilities. Let $\text{Cay}(G, S)$ denote the Cayley graph of G with respect to generating set S . Suppose $G = G_1 *_H G_2$ and $[G_i : H] \geq 2$, for $i = 1, 2$. If S_i is a set of generators for G_i then the graph $X = \text{Cay}(G, S_1 \cup H \cup S_2)$ has more than one end. Moreover, if $[G_1 : H] = [G_2 : H] = 2$ then X has two ends, otherwise X has infinitely many ends. If G is an HNN extension $G = G_1 *_H$, so that G_1 is a subgroup of G with isomorphic

finite subgroups H and $t^{-1}Ht$, then the Cayley graph $X = \text{Cay}(G, S_1 \cup \{t\})$ has more than one end.

If G has a Cayley graph X with more than one end then the cut system in Example 2.1 is not trivial and we can apply Theorem 7.2 to get a group action of G on the tree T . Then G splits over stabilizers of elements of a cut system (i.e. stabilizers of the edges of T). The splitting is non trivial as the graph X is vertex transitive and removing any separator in the cut system will leave at least two infinite components. The stabilizers of a cut A is finite, since it is a subgroup of the stabilizer of the finite set NA and the action of G on X is free. \square

Our results also provide a generalization of Stallings' Theorem to cases when the splitting group is not finite. Thus if a G -graph X with $G \backslash X$ finite has more than one vertex end then G has a non-trivial action on a tree T_n , which is a structure tree corresponding to the vertex end cut system of Example 2.1. The quotient $G \backslash T_n$ will be finite if $G \backslash X$ is finite and as T_n has infinite diameter, at least one orbit of edges of T_n will give a splitting of G . The splitting group will be the stabilizer of an edge in the selected orbit. This edge is a cut E in the cut system. Its stabilizer G_E will also stabilize NE . Since NE is finite, G_E will contain a subgroup of finite index which fixes each vertex in NE . Thus G_E has a subgroup of finite index which fixes a vertex of X . This subgroup may well be a proper subgroup of the stabilizer of this vertex. We summarize this in a theorem.

Theorem 8.4. *Let X be a G -graph with more than one vertex end and suppose $G \backslash X$ is finite, then G splits over a subgroup H such that a finite index subgroup of H fixes a vertex of X .*

Another possible application of vertex cuts is to the Kropholler Conjecture [18]. This arose out of work of Kropholler [17] on algebraic versions of the torus theorem for 3-manifolds.

Let H be a subgroup of the finitely generated group G . A subset A of a G -set is called H -finite if it is contained in the union of finitely many H -orbits, otherwise A is called H -infinite. We regard G as a G -set via the action of G on the left.

Conjecture 8.5 (Kropholler). *Let A be a subset of a finitely generated group G and let H be a subgroup of G such that $AH = A$. Let A and $G \backslash A$ be H -infinite and let $Ag \backslash A$ be H -finite, for all $g \in G$. Then G admits a non-trivial splitting over a group which is commensurable with a subgroup of H .*

If one could construct a G -graph X in which VX is the set of left cosets of H , which has more than one vertex end, then this conjecture would follow from the last theorem. One can get quite a long way in this direction. There will be a graph X in which VX is the set of left cosets of H , and such that $G \backslash X$ is finite. The set A will then determine a set E of vertices of this graph. The set NE is contained in finitely many H -orbits and since H fixes a vertex of X , NE has finite diameter in X . Both E and E^* have infinite diameter. This implies that both E and E^* contain rays. For more details we refer to [15, Theorem 3.5]. We have not been able to show that such a graph X can be constructed in which NE is finite, rather than just of finite diameter. If G is the commensurizer of H then one can construct X so that it is locally finite. A subset of VX will then be finite if and only if it has finite diameter. Thus the conjecture is true in this case. This was well known [6].

REFERENCES

- [1] D.E.Cohen. *Groups of cohomological dimension one* Springer Lecture Notes 245, 1972.
- [2] W. Dicks and M.J. Dunwoody. *Groups acting on graphs*. Cambridge University Press, Cambridge, 1989. Errata <http://mat.uab.es/dicks/>
- [3] C. Droms, B Servatius, H.Servatius. *The structure of locally finite two-connected graphs*. Electronic J. of Comb. **118** no. 17 (1995)
- [4] M.J. Dunwoody. Cutting up graphs. *Combinatorica* **2** 15-23 (1982)
- [5] M.J. Dunwoody. Inaccessible groups and protrees *J. Pure Applied Al.* **88** 63-78 (1993)
- [6] M. J. Dunwoody and E. L. Swenson, *The algebraic torus theorem*, Invent. Math. **140**, 605-637 (2000).
- [7] M.J. Dunwoody. Planar graphs and covers. Pre-print (2007), <http://www.personal.soton.ac.uk/mjd7/Pub.html>, arXiv:0708.0920.
- [8] A. Evangelidou and P. Papasoglu. Cuts and cyclic subsets of graphs. Preprint 2008.
- [9] H. Freudenthal. Über die Enden topologischer Räume und Gruppen. *Math. Zeitschr.* **33**, 692–713 (1931)
- [10] H. Freudenthal. Neuaufbau der Endentheorie. *Ann. of Math. (2)* **43**, 261–279, (1942)
- [11] H. Freudenthal. Über die Enden diskreter Räume und Gruppen. *Comm. Math. Helv.* **17**, 1–38 (1944)
- [12] H. A. Jung, M. E. Watkins. On the connectivities of finite and infinite graphs. *Monatsh. Math.* **83** no. 2, 121–131 (1977)
- [13] H. A. Jung, M. E. Watkins. Finite separating sets in locally finite graphs. *J. Combin. Theory Ser. B* **59**, 15–25 (1989)
- [14] B. Krön. End compactifications in non-locally finite graphs. *Proc. Phil Cambridge* **131**, 427-443 (2001).
- [15] B. Krön, R.G. Möller. Metric ends, fibers and automorphisms of graphs. *Math. Nachr.* **281** no. 1, 62–74 (2008).
- [16] B. Krön. Cutting up graphs revisited. Pre-print 2009.
- [17] P.H. Kropholler. An analogue of the torus decomposition theorem for certain Poincaré duality groups. *Proc. London Math. Soc.* **60**, 503-529 (1990).
- [18] G. Niblo and M. Sageev. The Kropholler conjecture. In *Guido's Book of Conjectures*, pages 106–107. 2006. Available at <http://www.fim.math.ethz.ch/preprints/mislinfinal.pdf>.
- [19] J.-P. Serre. *Trees*. Translated from the French original by John Stillwell. Corrected 2nd printing of the 1980 English translation. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003.
- [20] J.-P. Serre. *Trees*. Translated from the French by John Stillwell. Springer-Verlag, Berlin-New York, 1980.
- [21] J.R. Stallings. On torsion-free groups with infinitely many ends. *Annals of Mathematics (2)* **88**, 312–334 (1968)
- [22] J.R. Stallings. Groups of cohomological dimension one. Applications of Categorical Algebra (Proc. Sympos. Pure Math., Vol. XVIII, New York, 1968) pp. 124–128 Amer. Math. Soc., Providence, R.I. (1970)
- [23] J.R. Stallings. Group theory and three-dimensional manifolds. A James K. Whittemore Lecture in Mathematics given at Yale University, 1969. Yale Mathematical Monographs, 4. Yale University Press, New Haven, Conn.-London, 1971.
- [24] C. Thomassen, W. Woess. Vertex-transitive graphs and accessibility, *J. Combin. Theory Ser. B* **58**, 248–268 (1993)
- [25] W.T. Tutte. *Graph Theory*. Cambridge University Press, Cambridge, 1984.