ON THE STRUCTURE OF VERTEX CUTS
SEPARATING THE ENDS OF A GRAPH

GARETH R. WILKES

ON THE STRUCTURE OF VERTEX CUTS
SEPARATING THE ENDS OF A GRAPH

GARETH R. WILKES

Dinic, Karzanov and Lomonosov showed that the minimal edge cuts of a finite graph have the structure of a cactus, a tree-like graph constructed from cycles. Evangelidou and Papasoglu extended this to minimal cuts separating the ends of an infinite graph. In this paper we show that minimal vertex cuts separating the ends of a graph can be encoded by a succulent, a mild generalisation of a cactus that is still tree-like. We go on to show that the earlier cactus results follow from our work.

1. Introduction and definitions

Lying on the boundaries of several topic areas, vertex and edge cuts of graphs have been considered by graph theorists, network theorists, topologists and geometric group theorists, and the study of their structure has led to applications ranging from algorithms to classical group theoretic propositions.

Vertex cut pairs of finite graphs were studied by Tutte [1984], who showed that a graph possessing such cuts can be modelled with a tree. This was extended to infinite, locally finite graphs in [Droms et al. 1995]. Dunwoody and Krön [2015] then extended this work to cuts of other cardinalities, using vertex cuts to associate structure trees to graphs in a more general context.

This process of finding trees associated to graphs gives a way into geometric group theory. If, for instance, we find a structure tree for the Cayley graph of a group, then in light of the work of Bass and Serre [Serre 1980], we can obtain information about the group from its action on the tree. An example is Stallings' theorem [1968] on the classification of groups with many ends. The work of Dunwoody and Krön [2015] and of Evangelidou and Papasoglu [2014] yields more proofs of Stallings' theorem along these lines.

Dinic, Karzanov and Lomonosov [Dinic et al. 1976] showed that minimal edge cuts of a finite graph have, in addition to a tree-like nature, the finer structure of a cactus graph. For a recent elementary proof, see [Fleiner and Frank 2009]. Evangelidou and Papasoglu [2014] extended this, encoding all minimal edge cuts

separating the ends of an infinite graph by a cactus. The important stages in these proofs involve showing that certain collections of "crossing" cuts have a circular structure. In this paper we switch our attention to vertex cuts, showing that we can encode all minimal vertex end cuts of a graph by a tree-like structure called a succulent, which is a mild generalisation of a cactus. A traditional cactus is composed of cycles joined together at vertices in a tree-like fashion. For our succulents, we also allow cycles to attach along a single edge, again in a tree-like way. Once again the key step is to show that crossing cuts have a cyclic nature.

We will also show how the earlier cactus theorems can be regarded as special cases of our work, and discuss an application to certain finite graphs.

Let $\Gamma = (V, E)$ be a connected graph. If $K \subseteq V$ is a set of vertices of the graph, we denote by $\Gamma - K$ the graph obtained from Γ by removing K and all edges incident to K. K is called a *vertex cut* if K is finite and $\Gamma - K$ is not connected. If $A, B \subseteq \Gamma$ then we say K *separates* A and B if any path joining a vertex of A to a vertex of B intersects K.

A *ray* of Γ is an infinite sequence of distinct consecutive vertices of Γ. We say that two rays r_1, r_2 are *equivalent* if for any vertex cut K, all vertices of $r_1 \cup r_2$ except finitely many are contained in the same component of $\Gamma - K$. The *ends* of Γ are equivalence classes of rays. If K is a vertex cut of Γ, we say K is an *end cut* if there are at least two components of $\Gamma - K$ which contain rays. We say that an end cut is a *mincut* if its cardinality is minimal amongst end cuts of Γ. A mincut is said to *separate* ends e_1, e_2 of a graph if there are rays r_1, r_2 representing e_1, e_2 respectively such that r_1, r_2 are contained in different components of $\Gamma - K$. A mincut gives a partition of the set $\mathcal{E}$ of ends of the graph. Two mincuts are called *equivalent* if they give the same partition of $\mathcal{E}$. We denote the equivalence class of a mincut K by $[K]$, and write $K \sim L$ if K, L are equivalent.

A *succulent* is a graph constructed from cycles by joining cycles together at vertices or at single edges, in a "tree-like" fashion. We give a more formal definition of this as Definition 8.1 below. An *end vertex* of a succulent is one incident to at most two edges. We now state the main theorem of the paper.

Theorem 8.2. *Let Γ be a connected graph such that there are vertex end cuts of Γ with finite cardinality. There is a succulent $\mathcal{S}$ with the following properties*:

(1) *There is a subset A of vertices of $\mathcal{S}$ called the* anchors *of $\mathcal{S}$. If two anchors are adjacent, one of them is an end vertex of the graph. Every vertex of $\mathcal{S}$ not in A is adjacent to an anchor. We define an anchor cut of $\mathcal{S}$ to be a vertex cut containing no anchors which separates some anchors of $\mathcal{S}$. We say anchor cuts are equivalent if they partition A in the same way.*

(2) *There is an onto map f from the ends of Γ to the union of the ends of $\mathcal{S}$ with the end vertices of $\mathcal{S}$ which are anchors.*

(3) *There is a bijective map g from equivalence classes of minimal end cuts of Γ to equivalence classes of minimal anchor cuts of $\mathcal{S}$ such that ends e_1, e_2 of Γ are separated by $[K]$ if and only if $f(e_1)$, $f(e_2)$ are separated by $g([K])$.*

(4) *Any automorphism of Γ induces an automorphism of $\mathcal{S}$.*

2. Preliminaries

Definition 2.1. Given a mincut K, we call a component of $\Gamma - K$ *proper* if it contains an end, and a *slice* if not. Given a set of vertices C, its *boundary* ∂C is the set of those vertices not in C but adjacent to a vertex of C; and $C^* = V(\Gamma) - (C \cup \partial C)$.

It will be convenient to assume our graph contains no slices. In the following lemmas we show that we can do this by replacing Γ with another graph $\hat{\Gamma}$ which has the same ends and cuts, but no slices. The results in this section are adapted for our needs from more general results proved by Dunwoody and Krön [2015].

Lemma 2.2. *Let K, L be mincuts and C, D proper components of $\Gamma - K$, $\Gamma - L$. Suppose that both $C \cap D$ and $C^* \cap D^*$ contain an end. Then $\partial(C \cap D)$, $\partial(C^* \cap D^*)$ are mincuts,*

$$\partial(C \cap D) = (C \cap L) \cup (K \cap L) \cup (K \cap D),$$
$$\partial(C^* \cap D^*) = (C^* \cap L) \cup (K \cap L) \cup (K \cap D^*),$$

and

$$|C \cap L| = |K \cap D^*|,$$
$$|D \cap K| = |L \cap C^*|.$$

Proof. The boundaries $\partial(C \cap D)$, $\partial(C^* \cap D^*)$ are certainly end cuts, with

$$\partial(C \cap D) \subseteq (C \cap L) \cup (K \cap L) \cup (K \cap D),$$
$$\partial(C^* \cap D^*) \subseteq (C^* \cap L) \cup (K \cap L) \cup (K \cap D^*).$$

Consider the following diagram (see Figure 1), where a, b, c, d, u denote the cardinalities of the indicated sets. Let n be the cardinality of a mincut.
Then

$$a + c + u = n,$$
$$b + d + u = n.$$

Since $(C \cap L) \cup (K \cap L) \cup (K \cap D)$ is an end cut, we have

$$d + a + u \geq n,$$

and similarly

$$b + c + u \geq n.$$

Summing these and comparing with the equalities above, we find them to be equalities; it follows that $a = b$, $c = d$. $\quad\square$

Figure 1. Diagram for Lemma 2.2.

An analogous result holds when $C^* \cap D$, $D^* \cap C$ both contain ends.

Lemma 2.3. *If C, D are proper components of cuts K, L then there is a proper component of $\Gamma - K$ containing $C^* \cap L$.*

Proof. Since K, L are end cuts, one of the pairs $\{C \cap D, C^* \cap D^*\}$, $\{C^* \cap D, C \cap D^*\}$ both contain ends. Let A be the appropriate one of $C^* \cap D^*$, $C^* \cap D$. Then using Lemma 2.2, ∂A is a mincut. Let E be a component of A containing an end; then $\partial E = \partial A$ is a mincut. Let C_0^* be the component of C^* containing E. By Lemma 2.2 every vertex $x \in C^* \cap L$ is adjacent to E, so $x \in C_0^*$ and $C^* \cap L \subseteq C_0^*$. $\square$

Lemma 2.4. *A slice component of a mincut has empty intersection with each mincut. Distinct slices are disjoint. If Q is a slice, then no pair of elements of ∂Q are separated by any mincut.*

Proof. Let Q_1 be a slice component of $\Gamma - K$ for a mincut K and let L be a mincut, with a proper component D of $\Gamma - L$. By Lemma 2.3, there is a proper component of $\Gamma - K$ containing $C^* \cap L$, and C, a proper component, contains $C \cap L$. Q is disjoint from both of these, so $Q_1 \cap L = \varnothing$.

Suppose Q_2 is a slice component of $\Gamma - L$. We have $\partial Q_2 \subseteq L$, $\partial Q_1 \subseteq K$, and hence $Q_1 \cap \partial Q_2$, $Q_2 \cap \partial Q_1$ are both empty. The components Q_1, Q_2 are connected, so this implies that they are disjoint or equal.

Finally suppose $x, y \in \partial Q$ for a slice component Q of $\Gamma - K$ and x, y are separated by a mincut L. The slice Q is connected so there is a path in Q from x to y, which must intersect L, but we have seen this is impossible. $\square$

We will now show how to replace Γ with another graph $\hat{\Gamma}$ which has the same ends and cuts, but no slices. The vertex set $\hat{V}$ of $\hat{\Gamma}$ consists of those vertices of Γ which are contained in no slice. Two vertices $u, v \in \hat{V}$ are joined by an edge in $\hat{\Gamma}$ if and only if they are joined by an edge in Γ or if u, v lie in the boundary of some slice of Γ.

Lemma 2.5. *The graph $\hat{\Gamma}$ is connected and the mincuts of $\hat{\Gamma}$ are the same as the mincuts of Γ. There are no slices in $\hat{\Gamma}$. The ends of $\hat{\Gamma}$ are in bijection with the ends of Γ.*

Proof. First we show that if K is a mincut and C is a proper component thereof, then $\partial \hat{C}$, the boundary of $\hat{C} = C \cap \hat{\Gamma}$ as a subset of $\hat{\Gamma}$, is equal to K.

Suppose there is $x \in \partial \hat{C} - K$. If $x \in C$ then $x \in \hat{C}$, so $x \in C^*$. Also $x \in \partial \hat{C}$ so there is $y \in \hat{C}$ adjacent to x in $\hat{\Gamma}$. Then there is an edge from x to y in $\hat{\Gamma}$, but not in Γ; so x, y lie in the boundary of some slice Q of Γ. By Lemma 2.4, $K \cap Q = \varnothing$. The slice Q is connected, and Q intersects C (at y), so $Q \subseteq C$. We then have a path from x to y in Γ which is contained in Q except for its endpoint x, which is a path from C^* to C not intersecting K, a contradiction.

Suppose there exists $x \in K - \partial \hat{C}$; x has a neighbour y in $C - \hat{C}$. Then y is contained in a slice component Q of $\Gamma - L$ for a mincut L. If $K = L$ then C, Q are disjoint; but $y \in C \cap Q$. So $K \neq L$ and since $Q \subseteq C$ (Q does not intersect K but does intersect C), there is $z \in C \cap \partial Q \subseteq C \cap L$. Now z is not in any slice, so $z \in \hat{C}$. Then x, z are adjacent in $\hat{\Gamma}$; but $z \in \hat{C}$, $x \notin \hat{C} \cup \partial \hat{C}$, a contradiction.

Let us discuss the ends of $\hat{\Gamma}$. By definition, slices contain no rays. Thus if r is any ray in Γ, we can form a new ray in $\hat{\Gamma}$ by deleting any vertices in a slice; the extra edges added in the construction of $\hat{\Gamma}$ will ensure that this is a bona fide ray. If two rays are separated by a (not necessarily minimal) end cut K in Γ, then the union of $K \cap V(\hat{\Gamma})$ with the boundaries of any slices intersecting K gives an end cut separating the images of the rays in $\hat{\Gamma}$. Similarly, if two rays in $\hat{\Gamma}$ are separated by an end cut K in $\hat{\Gamma}$, then taking the union of K with any slice boundaries intersecting K gives an end cut separating the same rays in Γ. It follows that the ends of $\hat{\Gamma}$ are in a natural bijection with those of Γ.

The end cuts of $\hat{\Gamma}$ inherited from mincuts of Γ are indeed the minimal end cuts of $\hat{\Gamma}$. Suppose K is an end cut of $\hat{\Gamma}$ which is not also an end cut of Γ. Then two proper components of $\hat{\Gamma} - K$ are connected in Γ. A path between them can only not intersect K if it passes through a slice Q; but points on the boundary of Q are connected in $\hat{\Gamma}$ so we get a path between the two components in $\hat{\Gamma}$ as well, a contradiction. So all mincuts of $\hat{\Gamma}$ are mincuts of Γ as well, and the notion of minimality carries over to $\hat{\Gamma}$ too.

Finally, there are no slices in $\hat{\Gamma}$. Let C be a component of $\hat{\Gamma} - K$ for a mincut K of $\hat{\Gamma}$ (equivalently of Γ). Let C' be the component of $\Gamma - K$ containing C. Now C' cannot be a slice as it intersects $V(\hat{\Gamma})$. So C' contains an end of Γ, whence from above so does C. So C is not a slice. $\qquad\square$

For the rest of the paper we replace Γ with $\hat{\Gamma}$. As we have seen, the ends and cuts of the two graphs are the same, and this is all the structure with which we are concerned, so we lose nothing by doing this. All components of a cut are now proper.

We now start to prove some basic properties of mincuts, putting restrictions on cuts which "interact" with each other in some sense, and showing that a mincut does not interact with any but finitely many other mincuts. We first define what it means for cuts to not interact with each other. We are still following Dunwoody and Krön [2015] here, with some minor modifications.

Definition 2.6. Two cuts K, L are called *nested* if there are components E, F of $\Gamma - K$, $\Gamma - L$ respectively with $E \subseteq F$ or $F \subseteq E$.

Note that if K, L are nested and not equal with say $E \subseteq F$ then all components of $\Gamma - L$ except F are contained in the same component of $\Gamma - K$. This follows since there is an element of L in E^*, and by minimality all components of $\Gamma - L$ except F are connected to this vertex by paths which do not intersect $F \cup L$, and hence do not intersect $E \cup K$. Note also that these components are still connected in $\Gamma - K$ by similar reasoning. Conversely, all components of $\Gamma - K$ except one are contained in F.

Definition 2.7. A mincut is called an *A-cut* if it is nested with all other mincuts. It is called a *B-cut* if it separates Γ into exactly two components.

Lemma 2.8. *A mincut is either an A-cut or a B-cut.*

Proof. Let K be a mincut which is not an A-cut. Then there is a mincut L with which K is not nested. Let C be a (proper) component of $\Gamma - K$ and D a (proper) component of $\Gamma - L$. By Lemma 2.3, there is a component C_0^* of C^* containing $C^* \cap L$. We wish to show this is the only component of C^*. If there is another one C_1^* then $C_1^* \cap L$ is empty; C_1^* is connected so $C_1^* \subseteq D$ or $C_1^* \subseteq D^*$. In the first case, K, L are nested; so the second one happens no matter which component C_1^* we choose. So $D \cap C^* \subseteq C_0^*$. Also, $K = \partial C_1^*$ by minimality, so $C_1^* \subseteq D^*$ implies $K \cap D = \varnothing$. Then $D \subseteq C$ or $D \subseteq C^*$ (whence $D \subseteq C_0^*$); in either case, K and L are nested. This is a contradiction, so K is a B-cut. □

We call a set S of vertices a *tight x-y-separator* if $\Gamma - S$ has two distinct components A, B which are adjacent to all elements of S, with $x \in A$, $y \in B$.

Lemma 2.9. *For each integer k and every pair x, y of vertices of a graph, there are only finitely many tight x-y-separators of order k.*

Proof. We proceed by induction. If we take a path from x to y, any tight x-y-separator of order 1 would have to be a vertex on this path, so there are only finitely many of these.

Suppose the lemma holds for all tight x-y-separators of order k in all connected graphs. Take a path π from x to y in a graph Γ and suppose there are infinitely many tight x-y-separators of order $k + 1 \geq 2$. Then there is a vertex $z \in \pi - \{x, y\}$ which is contained in infinitely many of these separators. If S_1, S_2 are distinct such tight x-y-separators of order $k + 1$ in Γ then $S_1 - \{z\}$, $S_2 - \{z\}$ are distinct tight x-y-separators of order k in $\Gamma - \{z\}$, so there are infinitely many of these, giving the required contradiction. □

Lemma 2.10. *A mincut is nested with all but finitely many mincuts.*

Proof. Suppose K is a mincut and L is a mincut not nested with K. By Lemma 2.8 both K, L are B-cuts, with components C_1, C_2 of $\Gamma - K$ and D_1, D_2 of $\Gamma - L$. If $L \cap C_1$ was empty then by connectedness $C_1 \subseteq D_1$ or $C_1 \subseteq D_2$, both of which would imply that K, L were nested. Similarly none of $C_2 \cap L$, $D_1 \cap K$, $D_2 \cap K$ are empty. Then L is a tight x-y-separator for some $x \in K \cap D_1$, $y \in K \cap D_2$. There are only finitely many such separators for each pair x, y and only finitely many elements of K, so only finitely many such L are possible. $\qquad\square$

3. Crossing cuts

The complexity in the structure of mincuts comes from so-called "crossing" cuts, which we now define.

Definition 3.1. Let K, L be mincuts. Let $\mathcal{E}$ be the set of ends of Γ, and let $\mathcal{E} = K^{(1)} \sqcup K^{(2)} \sqcup \cdots \sqcup K^{(r)}$, $\mathcal{E} = L^{(1)} \sqcup L^{(2)} \sqcup \cdots \sqcup L^{(s)}$ be the partitions of $\mathcal{E}$ given by K, L respectively. We say $[K]$, $[L]$ *cross* if, possibly after relabelling, $K^{(i)} \cap L^{(j)} \neq \varnothing$ for i, $j = 1, 2$. We write $K + L$.

The following is a direct consequence of Lemma 2.8, having removed slices from our graph.

Lemma 3.2. *If $[K]$, $[L]$ cross then $\Gamma - K$, $\Gamma - L$ have exactly two components.*

Later we will show that crossing mincuts possess a cyclic structure. Initially, however, we shall just consider two or three crossing cuts.

Lemma 3.3. *Let $[K]$, $[L]$ be crossing classes of mincuts. Let $\Gamma - K = C_1 \sqcup C_2$, $\Gamma - L = D_1 \sqcup D_2$. Then $|C_1 \cap L| = |C_2 \cap L| = |D_1 \cap K| = |D_2 \cap K|$; i.e., $K \cup L$ splits into four equal pieces, plus the "centre" $U = K \cap L$.*

Proof. This follows from two applications of Lemma 2.2. $\qquad\square$

In the case of edge cuts, one can also show that the centre $K \cap L$ is empty, but in the case of vertex cuts this fails to be true. As we will show in Lemma 3.5 below, the centre is in some sense distinguished, but this result must wait until we have placed some restrictions on the division of a graph produced by three cuts.

Let K, L, M be mincuts with K crossing L and L crossing M, and let $C_1 \sqcup C_2$, $D_1 \sqcup D_2$, $E_1 \sqcup E_2$ be the components of $\Gamma - K$, $\Gamma - L$, $\Gamma - M$ respectively. A priori, these three cuts could divide Γ into eight components each containing an end. The natural diagram with which to illustrate this would be a suitably divided cube. To produce this in 2D we divide the cube into three slices as shown in Figure 2.

We now rule out certain arrangements of ends of the graph.

Lemma 3.4. *It is not possible for each of $C_1 \cap D_1 \cap E_1$, $C_2 \cap D_1 \cap E_2$, $C_2 \cap D_2 \cap E_1$, $C_1 \cap D_2 \cap E_2$ to contain an end (or any arrangement obtained from this by relabellings).*

Figure 2. The notation shows that, for example, (1) is $C_1 \cap D_1 \cap E_2$ and (2) is $C_2 \cap D_2 \cap M$.

Proof. Denote by $a, \ldots, u$ the cardinalities of the various subgraphs as shown in Figure 3; the ϵ_i indicate the presence of ends.

Let n be the cardinality of a mincut. Then $|K| = |L| = |M| = n$, so

$$n = a + c + l + j + m + p + r + t + u,$$
$$n = e + f + g + h + q + r + s + t + u,$$
$$n = b + d + i + k + m + p + q + s + u.$$

We also have an end cut separating each ϵ_i from the others; these yield, in order,

$$n \le a + b + e + m + q + t + u,$$
$$n \le c + d + g + m + r + s + u,$$
$$n \le k + l + h + p + s + t + u,$$
$$n \le i + j + f + p + q + r + u.$$

Summing these four, we have

$$4n \le (a + c + l + j + m + p + r + t + u) + (e + f + g + h + q + r + s + t + u)$$
$$+ (b + d + i + k + m + p + q + s + u) + u$$

$$= 3n + u,$$

whence $u = n$, everything else vanishes, and $K = L = M$, each separating the graph into at least four components, contradicting Lemma 3.4. $\square$

Figure 3. Diagram for Lemma 3.4.

Figure 4. Diagram for Lemma 3.5.

Note that this result implies that the three cuts split the graph into at most six components containing ends. Since K crosses L, there are at least four such components. A quick exercise in filling in corners with ends subject to the crossings and Lemma 3.4 shows that, following relabellings, $C_1 \cap D_1 \cap E_1$, $C_2 \cap D_1 \cap E_2$, $C_1 \cap D_2 \cap E_1$, $C_2 \cap D_2 \cap E_2$ all contain ends (with possibly other corners also).

Lemma 3.5. *Let K, L, M be mincuts with K crossing L and L crossing M (in particular, if K is equivalent to M). Then $K \cap L = L \cap M$, and $L \cap C_1 = L \cap E_1$, $L \cap C_2 = L \cap E_2$.*

Proof. Retain the notations of the previous lemma and see Figure 4. Again since K, L, M are mincuts,

$$n = a + c + l + j + m + p + r + t + u,$$
$$n = e + f + g + h + q + r + s + t + u,$$
$$n = b + d + i + k + m + p + q + s + u,$$

and again, considering end cuts separating a corner containing an end ϵ_i from the others, we have

$$n \le a + b + e + m + q + t + u,$$
$$n \le c + d + g + m + r + s + u,$$
$$n \le i + l + e + p + q + t + u,$$
$$n \le j + k + g + p + r + s + u.$$

Summing these,

$$4n \le (a + c + l + j + m + p + r + t + u) + (b + d + i + k + m + p + q + s + u)$$
$$+ 2e + 2g + q + r + s + t + 2u$$
$$= 2(e + f + g + h + q + r + s + t + u) + 2n - 2f - 2h - q - r - s - t$$
$$= 4n - (2f + 2h + q + r + s + t),$$

whence $f = h = q = r = s = t = 0$, so that $K \cap L = L \cap M$.   □

Lemma 3.6. *A cut is crossed by at most finitely many cuts.*

Proof. If two cuts cross they are not nested, so this lemma follows directly from Lemma 2.10. □

4. Half-cuts

It follows from the last section's results that each mincut in a crossing system can be decomposed into three pieces; two "half-cuts" and a "centre". We now prove some facts about these half-cuts, which enable us to arrange the half-cuts on a circle.

Definition 4.1. If K, M are mincuts (more properly, classes of mincuts under $\sim$, but we will often pass over this technicality), we write $K \# L$ if there are mincuts $K = L_0, L_1, \ldots, L_n = M$ such that $L_0 + L_1 + \cdots + L_n$; that is, L_0 crosses L_1, L_1 crosses L_2 and so on. L_0 may or may not cross L_2. Note that $\#$ is an equivalence relation on $\sim$-classes of mincuts, decomposing these into equivalence classes, which we call #-classes.

By Lemma 3.5, elements K of a #-class have a unique decomposition $K = K_1 \cup U \cup K_2$, where if $K + L$ then $K \cap L = U$ and K_1, K_2 are in different components of $\Gamma - L$. From the same lemma, this U is the same for all cuts in the #-class; we call it the *centre* of the #-class. Also, $|K_1| = |K_2|$ and this cardinality is again the same across the class. The K_i are called the *half-cuts* of the #-class. We now prove a series of lemmas clarifying the structure of a #-class and its half-cuts.

Lemma 4.2. *If K, M are mincuts in the same #-class then either $K + M$ or there is an L in this #-class such that $K + L + M$; that is, $K + L$ and $L + M$.*

Proof. By definition we have a sequence of cuts $K = L_0, L_1, \ldots, L_n = M$ such that $L_0 + L_1 + \cdots + L_n$. Take a shortest such sequence, and suppose $n \geq 3$. We will show we can find a shorter sequence. Without loss of generality we can assume $n = 3$. Let $\mathcal{E} = L_i^{(1)} \cup L_i^{(2)}$ be the partition induced by L_i. The fact that K does not cross L_2, and similar facts, give us, after relabelling,

$$K^{(1)} \subseteq L_2^{(1)}, \quad L_2^{(2)} \subseteq K^{(2)},$$
$$M^{(2)} \subseteq L_1^{(2)}, \quad L_1^{(1)} \subseteq M^{(1)},$$

whence the crossings give us that each of $M^{(2)} \cap K^{(2)}$, $M^{(1)} \cap K^{(2)}$, $M^{(1)} \cap K^{(1)}$ is nonempty. Hence $K + M$ unless $M^{(2)} \cap K^{(1)}$ is empty, and hence $K^{(1)} \subseteq M^{(1)}$. It is this that allows us to place the ends ϵ_3, ϵ_6 in Figure 5, and hence to conclude that $K + (L_{11} \cup U \cup L_{22}) + M$, where for instance L_{11} is the half-cut of L_1 lying in the $K^{(1)}$-component of $\Gamma - K$. □

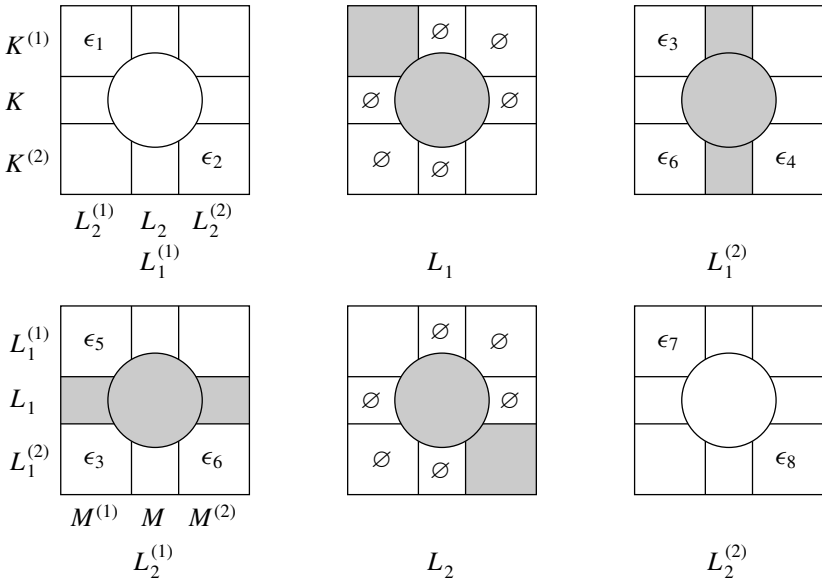**Corollary 4.3.** *There are only finitely many cuts in a #-class and hence only finitely many half-cuts.*

Figure 5. Diagram for Lemma 4.2. The cut $L_{11} \cup U \cup L_{22}$ is shown shaded in both diagrams.

Lemma 4.4. *Let K_1, M_1 be half-cuts in the same #-class. Then $K_1 \cup U \cup M_1$ is a mincut if and only if there are mincuts K', M' containing K_1, M_1 as half-cuts respectively such that $K' + M'$.*

Proof. One direction is clear. For the other, pick K_2, M_2 such that $K = K_1 \cup U \cup K_2$, $M = M_1 \cup U \cup M_2$ are cuts in this #-class. Then either $K + M$, in which case we are done, or there is L such that $K + L + M$. Now $K_1 \cup U \cup M_1$ is a cut; hence we have ϵ_5 in Figure 6. Then $K_1 \cup U \cup M_2$, $K_2 \cup U \cup M_1$ cross. $\square$

Definition 4.5. Two half-cuts K_1, L_1 in the same #-class are *equivalent* if whenever K_2 is a half-cut such that $K_1 \cup U \cup K_2$ is a mincut, then $L_1 \cup U \cup K_2$ is an equivalent cut and vice versa.

Figure 6. Diagram for Lemma 4.4. The cut $K_1 \cup U \cup M_2$ is shown shaded.

Figure 7. Diagram for Lemma 4.6.

Two half-cuts K_1, L_1 in the same #-class are *quasiequivalent* if there is a half-cut K_2 such that $K_1 \cup U \cup K_2$ is a mincut and $L_1 \cup U \cup K_2$ is an equivalent cut.

Lemma 4.6. *If two half-cuts K_1, M_1 form a cut then they are not quasiequivalent.*

Proof. Let $K = K_1 \cup U \cup M_1$ be the cut formed by hypothesis. Let L_1 be some other half-cut; we will show that $K_1 \cup U \cup L_1$ is not equivalent to $L_1 \cup U \cup M_1$ as cuts, and hence that K_1, M_1 are not quasiequivalent. Let L_2 be a half-cut such that $L = L_1 \cup U \cup L_2$ is in the #-class. If $L + K$ then the result is clear. If not, there is a mincut N such that $K + N + L$; without loss of generality take K_1, L_1 to be in the same component of $\Gamma - N$. Then $L_1 \cup U \cup M_1$ is a cut, and from Figure 7 we see that either $K_1 \cup U \cup L_1$ is not an end cut or it is not equivalent to $L_1 \cup U \cup M_1$. $\square$

Lemma 4.7. *Let $K = K_1 \cup U \cup K_2$ be a cut in the #-class and let M_1 be a half-cut in the same class not quasiequivalent to either K_1, K_2. Then there is M_2 such that $M_1 \cup U \cup M_2$ is a cut crossing K. Hence $K_1 \cup U \cup M_1$ and $K_2 \cup U \cup M_2$ are cuts.*

Proof. Let M_2 be a half-cut such that $M = M_1 \cup U \cup M_2$ is a cut of the class; see Figure 8. If $K + M$, we are done. Otherwise, there is a cut L with $K + L + M$. After possibly relabelling the K_i, we can assume that K_1, M_1 are in the same component of $\Gamma - L$. If there is an end in $C_1 \cap D_1 \cap E_2$ then $M_1 \cup U \cup L_2$ is a cut crossing K. If there is an end in $C_2 \cap D_1 \cap E_1$ then $M_1 \cup U \cup L_1$ is a cut crossing K. If neither of these happens, then $K_2 \cup U \cup M_1$ is equivalent to $K_2 \cup U \cup K_1$, a contradiction


**Figure 8.** Diagram for Lemma 4.7.

Figure 9. Diagram for Lemma 4.8.

(we note that these cuts are genuinely equivalent, since the presence of "links" such as L_1 guarantees that ends which appear to be connected up really are). $\square$

Lemma 4.8. *Quasiequivalence is an equivalence relation. If K_1, L_1 are quasi-equivalent and L_2 is a half-cut such that $L_1 \cup U \cup L_2$ is in the #-class then $K_1 \cup U \cup L_2 \sim L_1 \cup U \cup L_2$.*

Proof. Let K_2 be such that $K = K_1 \cup U \cup K_2$ is in the #-class; see Figure 9. The cut K does not cross $L = L_1 \cup U \cup L_2$ since in this case $K_1 \cup U \cup L_1$ would be a cut, so K_1, L_1 are not quasiequivalent by Lemma 4.6, giving a contradiction. Then there is an N such that $K + N + L$. Again, $K_1 \cup U \cup L_1$ is not a cut, so there are no ends in certain corners as indicated. Then $K_1 \cup U \cup L_2 \sim L_1 \cup U \cup L_2$ as required, noting again that ends which appear connected actually are so that the cuts are genuinely equivalent.

As for quasiequivalence being an equivalence relation, it is clearly symmetric and reflexive. If M_1 is another half-cut quasiequivalent to L_1, then by the above

$$K_1 \cup U \cup L_2 \sim L_1 \cup U \cup L_2 \sim M_1 \cup U \cup L_2,$$

so K_1, M_1 are quasiequivalent. $\square$

Lemma 4.9. *Let $K = K_1 \cup U \cup K_2$ be a cut in the #-class and let L_1, M_1 be half-cuts in the same class not quasiequivalent to K_1, K_2. Then either $L_1 \cup U \cup M_1$ is a cut crossing K or L_1, M_1 are contained in the same component of $\Gamma - K$.*

Proof. By Lemma 4.7, we can complete L_1 to a cut crossing K, so that L_1 separates some ends of a component of $\Gamma - K$, and it works similarly for M_1. If the two half-cuts are in different components, then $L_1 \cup U \cup M_1$ is a cut crossing K provided it separates $\Gamma - K$ into two components, as indeed it must. $\square$

5. Separation systems

We now turn our attention to demonstrating that the half-cuts of a system have a cyclic structure. We will do this by showing that they satisfy a certain axiomatic

system, which implies that they can be arranged cyclically in a fashion compatible with their cut structure. This axiomatic structure is taken from [Huntington 1935].

Definition 5.1. A separation relation on a set Z is a relation $R \subseteq Z^4$ satisfying the following axioms. We write $abcd$ if $(a, b, c, d) \in R$.

(1) If $abcd$ then a, b, c, d are distinct.

(2) There are a, b, c, d such that $abcd$, i.e., $R \neq \varnothing$.

(3) If $abcd$ then $bcda$.

(4) If $abcd$ then $\neg(abdc)$.

(5) There are a, b, c, d such that $abcd$ and $dcba$.

(6) If $abcd$ and $x \in Z$ is another element then either $axcd$ or $abcx$.

Lemma 5.2. *Let Z be a set equipped with a separation relation. Then*

(1) *if $a, b, c, d \in Z$ are distinct, then at least one of the twenty-four tetrads $abcd, abdc, \dots, dcba$ is true,*

(2) *if $abcd$ then $dcba$,*

(3) *if $abxc$ and $abcy$ then $abxy$,*

(4) *if $abcx$ and $abcy$ then $abxy$ or $abyx$,*

(5) *if $abcx$ and $abcy$ then $acxy$ or $acyx$,*

where in the last three statements distinct letters are assumed to represent different elements of Z.

Proofs can be found in [Huntington and Rosinger 1932] along with further similar propositions.

Lemma 5.3. *Let Z be a finite set with a separation relation. For each z, there are unique a, b such that for all $c \in Z - \{z, a, b\}$, we have $azbc$. We call these the elements adjacent to z.*

Proof. We approach existence by induction. For $|Z| = 4$, the result is trivial. Assume it is true for all separation relations with $|Z| = n$, and suppose $|Z| = n + 1$. Remove an element d of Z not equal to z to leave a smaller separation relation, and let a, b be the elements adjacent to z in this new relation, so that for all $c \in Z - \{z, a, b, d\}$, we have $azbc$.

By Lemma 5.2, one of $azbd$, $adzb$, $azdb$ holds. If $azbd$ holds then a, b are adjacent to z in Z. If not, without loss of generality, $azdb$. We claim a, d are adjacent to z in Z. By Lemma 5.2 above, if $c \in Z - \{z, a, d, b\}$ then $azdb$ and $azbc$ imply $azdc$.

For the uniqueness part, suppose there are two such pairs a_1, b_1, a_2, b_2. If any of these coincide we have an immediate contradiction to part (4) of the definition of the relation. So suppose they are all distinct. Then $a_1 z b_1 a_2$, $a_1 z b_1 b_2$ imply $a_1 z a_2 b_2$ or $a_1 z b_2 a_2$, both of which contradict $a_2 z b_2 a_1$.   □

Lemma 5.4. *Let Z be a finite set equipped with a separation relation. Then there is a map $F : Z \to S^1$ such that for $a, b, c, d \in Z$, $abcd$ if and only if $F(b)$ and $F(d)$ lie in different components of $S^1 - \{F(a), F(c)\}$; i.e., Z is isomorphic to a finite subset of the circle under its natural separation relation.*

Proof. We will proceed by induction. Pick an element $z \in Z$ and take a separation-preserving map $\overline{F} : Z - \{z\} \to S^1$. By the previous lemma, there are elements a, b of Z adjacent to z. We will map z to the circle by placing it between $\overline{F}(a), \overline{F}(b)$, but first we must show these are adjacent on the circle. If not, there are c, d so that $\overline{F}(a)\overline{F}(c)\overline{F}(b)\overline{F}(d)$, whence $acbd$. But $azbc$, $azbd$ imply $abcd$ or $abdc$, both contradicting $acbd$. So $\overline{F}(a), \overline{F}(b)$ are adjacent on the circle, and we can define $F : Z \to S^1$ by setting $F = \overline{F}$ on $Z - \{z\}$ and $F(z)$ to lie between $\overline{F}(a), \overline{F}(b)$ on the circle.

A full proof that this F works would be lengthy and uninformative, so we just indicate the main steps; the remainder is just use of axioms and Lemma 5.2. We inherit from $\overline{F}$ that any relations not involving z are preserved. Let $zABC$ be another relation and suppose A, B, C are distinct from a, b; the other cases are easier. Then we have $azbA$, $azbB$, $azbC$, $zABC$ from which we deduce $aABC$, $bABC$. These relations carry over to the circle under F, as do $azbA$, $azbB$, $azbC$ by construction. From these relations on the circle, we then find $F(z)F(A)F(B)F(C)$. $\square$

Definition 5.5. Let Z be the set of quasiequivalence classes of half-cuts of a #-class. We define a separation relation R on Z by setting $(a, b, c, d) \in R$ if and only if $ac + bd$, where ac denotes the cut $K_1 \cup U \cup L_1$ and K_1, L_1 are representatives of a, c and so on.

Lemma 5.6. *This is well-defined; i.e., it does not matter which representatives of quasiequivalence classes we choose. Furthermore, it is a bona fide separation relation.*

Proof. Well-definedness follows immediately from Lemma 4.8. Parts (1)–(5) of the definition of a separation relation are trivial. For part (6), by Lemma 4.9 either $abcx$ or b, x are in the same component of $\Gamma - ac$; and either $axcd$ or x, d are in the same component of $\Gamma - ac$. But b, d are in different components of $\Gamma - ac$ so one of $abcx$, $axcd$ holds. $\square$

Hence we have:

Proposition 5.7. *To each #-class we can associate a cycle where each vertex represents a quasiequivalence class and each cut of the #-class is associated to a vertex cut of the cycle, with the notions of crossing preserved.*

6. The structure of a #-class

We are now in a position to characterize the structure of a #-class. Let $[K_1]_q$ denote the quasiequivalence class of a half-cut K_1. From Proposition 5.7, there are two quasiequivalence classes adjacent to $[K_1]_q$ in this #-class. If L_1 is a half-cut in the #-class not in $[K_1]_q$ or either quasiequivalence class adjacent to it, then Lemma 4.8 implies that

$$K_1 \cup U \cup L_1 \sim K_1' \cup U \cup L_1$$

for all $K_1' \in [K_1]_q$. So only the two quasiequivalence classes adjacent to $[K_1]_q$ can contain L_1 such that $K = K_1 \cup U \cup L_1$ and $K' = K_1' \cup U \cup L_1$ are not equivalent for $K_1' \in [K_1]_q$.

   How can these cuts be nonequivalent? We recall that by minimality every component left by a mincut is connected to every element of that cut. Thus in the "larger component" left by the cut, i.e, the one containing half-cuts in the same class, every vertex is connected to the half-cuts in this "component", which is thus genuinely connected. Thus one part of the partitions induced by K, K' is the same. The others can only differ if at least one of the cuts splits Γ into more than two parts, and hence splits the "smaller" component into more than one part. Suppose K intersects one of the smaller components of $\Gamma - K'$. Then each end not in the larger component of $\Gamma - K'$ is connected to each vertex of the part of K in the smaller component; hence K' splits Γ into exactly two components. If K does not intersect one of the smaller components of $\Gamma - K'$, then since $K \neq K'$ and K, K' have the same cardinality, K' intersects one of the smaller components of $\Gamma - K$, whence K splits Γ into exactly two components.

   Hence, having chosen L_1, there are at most two equivalence classes of cuts formed from L_1 and $[K_1]_q$. By symmetry, there are at most two equivalence classes of cuts formed from K_1 and $[L_1]_q$. From these discussions, it follows that for each quasiequivalence class adjacent to $[K_1]_q$, there are at most two equivalence classes of cuts formed by these two classes; one producing a split of Γ into two components, the other more. Hence there are at most four equivalence classes of half-cuts within $[K_1]_q$.

   We now define the structure by which we model the #-class. For edge cuts this would be a simple cycle, but here we need extra complexity to deal with the possibility of splitting the graph into more than two components.

Definition 6.1. A *ring* is constructed as follows. Take a finite cycle of vertices and attach to each edge some number of triangles by identifying an edge of the triangle with the edge of the cycle. The vertex of the triangle not included in the original cycle is called an *anchor*; see Figure 11.

Figure 10. A 3-vertex connected to two 1-vertices, and the schematic representation of this.

Definition 6.2. An *n-vertex* will be a copy of the complete graph on *n* vertices; we will say it is connected to a vertex if there is an edge from the vertex to each constituent vertex of the *n*-vertex. We will depict a 3-vertex as a triangle and only draw one edge from it to each vertex to which it is connected; see Figure 10.

We now associate to each #-class an appropriate ring encoding the cuts formed by half-cuts in the class. First use Proposition 5.7 to form a cycle with one vertex for each quasiequivalence class. For each pair of adjacent quasiequivalence classes, find half-cuts in those classes separating Γ into as many components as possible, and attach one fewer anchors than this between the two classes in the cycle (one fewer to account for the "large" component). If a quasiequivalence class contains more than one equivalence class, insert an extra vertex into the cycle here. If we

Figure 11. A ring, with the anchors replaced by 3-vertices.

Figure 12. A #-class and its associated ring. Hexagons represent 6-vertices and arrows ends.

"thicken" up the anchors to 3-vertices to remove cut-points, there is now a bijective correspondence between equivalence classes of cuts formed from half-cuts of the #-class and equivalence classes of cuts of the ring, where we treat the anchors as ends for the purpose of equivalence etc.; see Figure 12.

7. Pretrees

We now proceed towards the central theorem of the paper. First we seek to impose a tree structure on the #-classes and the other cuts; then we will reintroduce the extra complexity. We will do this using pretrees, which we now define.

Definition 7.1. Let $\mathcal{P}$ be a set and let $R \subseteq \mathcal{P}^3$ be a ternary relation on $\mathcal{P}$. If $(x, y, z) \in R$ then we write xyz and say y is between x, z. A set $\mathcal{P}$ equipped with this relation is a *pretree* if the following hold:

(1) If xyz then $y \neq x, z$, and there are no x, y such that xyx.

(2) If xyz then zyx.

(3) For all x, y, z, if xyz then $\neg(xzy)$.

(4) If xzy and $w \neq x, y, z$ then xzw or yzw.

If there is no z such that xzy we say x, y are adjacent.

A pretree is called *discrete* if for any $x, y \in \mathcal{P}$ there are at most finitely many z such that xzy.

It should perhaps be noted that despite us using the word "between" this is not a betweenness relation in the usual sense of the word as, for example, in [Huntington 1935]. Let $\mathcal{P}$ be a discrete pretree. We will describe briefly how to pass from $\mathcal{P}$ to a tree; a fuller description may be found in [Bowditch 1999].

We call a subset H of $\mathcal{P}$ a *star* if all $a, b \in H$ are adjacent. We now define a tree T as follows:

$$V(T) = \mathcal{P} \cup \{\text{maximal stars of } \mathcal{P}\},$$

$$E(T) = \{(v, H) : v \in \mathcal{P}, v \in H, H \text{ a maximal star}\}.$$

We show that T is indeed a tree. If $x, y \in \mathcal{P}$ then by discreteness there are only finitely many z between x, y. From among these z we can then find $z_1, \ldots, z_n$ such that x is adjacent to z_1, z_i is adjacent to z_{i+1} and z_n is adjacent to y, giving a path in T from x to y. Hence T is connected.

If T contains a circuit then there are $x_1, \ldots, x_n$ in $\mathcal{P}$ such that x_i is adjacent to x_{i+1} but not to x_{i+2} for each $i \in \mathbb{Z}_n$. Then there is y such that $x_i y x_{i+2}$. If $y \neq x_{n+1}$ then either $x_i y x_{i+1}$ or $x_{i+1} y x_{i+2}$, both of which are forbidden. So $x_i x_{i+1} x_{i+2}$. We claim $x_1 x_i x_{i+1}$ holds for all $i \leq n$ by induction. Since $x_1 x_{i-1} x_i$ and $x_{i-1} \neq x_{i+1}$, either $x_1 x_{i-1} x_{i+1}$ holds or we have a contradiction. Since $x_{i-1} x_i x_{i+1}$ and $x_1 \neq x_i$,

Figure 13. Diagram for Lemma 7.3.

either $x_1 x_i x_{i+1}$ or $x_{i-1} x_i x_1$; so to avoid contradiction, $x_1 x_i x_{i+1}$. But then we have $x_1 x_{n-1} x_n$; but x_1, x_n were supposed to be adjacent. The contradiction means T is a tree.

We now prove some lemmas which will allow us to define a pretree of cut classes.

Definition 7.2. We call a mincut *isolated* if it does not cross any mincut, and hence is not contained in any #-class.

A cut is a *corner cut* of a #-class if it is (equivalent to) a cut formed from two half-cuts of the class but is not itself in the class. We call a mincut *totally isolated* if it does not cross any mincut and is not a corner cut of any #-class.

Lemma 7.3. *Corner cuts are isolated.*

Proof. Let Q be a #-class and let $K = K_1 \cup U \cup K_2$ be a corner cut of Q. Suppose there is a cut L with $K + L$. Then L separates some ends of each component of $\Gamma - K$. Let M_1, M_2 be half-cuts in Q adjacent to K, with K_1 adjacent to M_2 and K_2 adjacent to M_1, with no quasiequivalences present; see Figure 13. Either L crosses $K_1 \cup U \cup M_1$ or all ends of the component of $\Gamma - K_1 \cup U \cup M_1$ containing M_2 are in the same component of $\Gamma - L$, whence L crosses $K_2 \cup U \cup M_2$. So L, hence K, are in the #-class Q, a contradiction. $\qquad\square$

Each #-class Q induces two partitions

$$\mathcal{E} = Q^{(1)} \sqcup \cdots \sqcup Q^{(m)},$$
$$\mathcal{E} = \bar{Q}^{(1)} \sqcup \cdots \sqcup \bar{Q}^{(m')}$$

of the ends of Γ. In one partition, which we call the fine partition and denote without bars, each member of the partition corresonds to one of the anchors in the ring representing Q; and for each $Q^{(i)}$, there is a corner cut of Q separating $Q^{(i)}$ from all the other $Q^{(j)}$. For the other partition, the coarse partition, we identify those $Q^{(i)}$ together which lie between the same two adjacent half-cuts. Then in the coarse partition we can distinguish between members $\bar{Q}^{(i)}$ using only cuts properly in the #-class Q; for the fine partition, we may need corner cuts also. We recall

also that a cut K also gives a partition of the ends of Γ:

$$\mathcal{E} = K^{(1)} \sqcup \cdots \sqcup K^{(n)}.$$

Lemma 7.4. *Given a cut K and a #-class Q, with K neither in Q nor a corner cut of it, there are i, j such that all $Q^{(k)}$ except $Q^{(i)}$ are contained in $K^{(j)}$; i.e.,*

$$\coprod_{k \neq i} Q^{(k)} \subseteq K^{(j)}$$

or

$$\coprod_{k \neq j} K^{(k)} \subseteq Q^{(i)}.$$

We say K divides $Q^{(i)}$.

Proof. Suppose K is an A-cut. Then it is nested with every cut and corner cut of Q, hence the result.

Otherwise K is a B-cut, separating Γ into two components. If the result is not true, then both $K^{(i)}$ intersect at least two $Q^{(i)}$.

Suppose a $Q^{(i)}$ intersects both $K^{(i)}$. Let M be the corner cut of Q splitting off $Q^{(i)}$. If M is a B-cut, then $K + M$, giving a contradiction. Otherwise, M is nested with K, whence a $K^{(i)}$ is contained in a $Q^{(i)}$, again giving a contradiction.

If both $K^{(i)}$ contain two $Q^{(i)}$ not between two adjacent half-cuts (in the ring representing Q), we can find a cut of Q crossing K. So for say $K^{(1)}$, all the $Q^{(i)}$ contained in $K^{(1)}$ lie between two adjacent half-cuts of Q. Let M be the corner cut corresponding to these half-cuts.

M is necessarily an A-cut, and hence is nested with K. As in the discussion of quasiequivalent cuts earlier, K can only intersect the "large" component of $\Gamma - M$, that containing the other half-cuts of Q; conversely M does not intersect the large component C_1 of $\Gamma - K$. Pick another half-cut L_1 in Q, with $L = M_1 \cup U \cup L_1$ a cut of Q. The cut L also only intersects the large component E_1 of $\Gamma - M$. With suitable labelling of the $L^{(i)}$, we have

$$L^{(1)} \subseteq M^{(1)}, \quad L^{(1)} \subseteq K^{(1)}, \quad K^{(2)} \subseteq L^{(2)}, \quad K^{(1)} = M^{(1)}.$$

Hence we have the arrangement shown in Figure 14.

Then, using the notations from Figure 14, we have

$$a + m + p + l + t + u = n,$$
$$d + m + p + k + s + u = n,$$
$$e + t + u + h + s = n,$$
$$a + m + e + t + u \geq n,$$
$$p + l + e + t + u \geq n,$$
$$p + k + u + s \geq n,$$

Figure 14. Diagram for Lemma 7.4.

where n is the cardinality of a mincut. Immediately $d = m = 0$. Furthermore, since L_2, M_1, M_2 are half-cuts in the same class, $s = e + t + h$. Then

$$
\begin{aligned}
2n &\leq (a + e + t + u) + (p + l + e + t + u) \\
&= (a + p + l + t + u) + (e + e + t + u) \\
&\leq n + e + t + h + s + u \\
&= 2n.
\end{aligned}
$$

Then all the inequalities are equalities; hence $t = h = 0$, and K decomposes into U together with two equal half-cuts; and choosing L_1 appropriately we find that these half-cuts are quasiequivalent to the half-cuts of M, so K was a corner cut of Q. $\square$

Lemma 7.5. *Given two #-classes Q, R, all cuts in R divide the same $Q^{(i)}$.*

Proof. Note first that the cuts in R do divide a $Q^{(i)}$ because they are not isolated, hence not corner cuts, and are not in Q. Suppose $K \in R$ divides $Q^{(i)}$ and $L \in R$ divides $Q^{(j)}$, with $i \neq j$; see Figure 15.

If there is one, take a cut M crossing K and L. We have $K^{(2)} \subseteq Q^{(i)}$ and $L^{(2)} \subseteq Q^{(j)}$, so M contradicts Lemma 7.4.

Then $K + L$. Take a cut $M \in Q$ separating $Q^{(i)}$ from $Q^{(j)}$, and let $N = K_2 \cup U \cup L_2$. The cut N separates some ends of $Q^{(i)}$ and $Q^{(j)}$; it is not nested with M, and hence is a B-cut and crosses M, giving a contradiction. $\square$

Figure 15. Diagram for Lemma 7.5. The cut N is shown shaded.

Lemma 7.6. *Given two totally isolated cuts K, L, we have that K divides only one $L^{(i)}$; i.e., there are i, j such that*

$$\bigsqcup_{k\neq i} L^{(k)} \subseteq K^{(j)}, \quad \bigsqcup_{k\neq j} K^{(k)} \subseteq L^{(i)}.$$

Proof. If K, L are nested, the result is immediate. If not, they are both B-cuts, and the result follows since they do not cross. □

We now define a pretree encoding the mincuts of Γ. Let $\mathcal{P}$ be the set of all #-classes of Γ and all equivalence classes of totally isolated cuts of Γ. Given $x, y, z \in \mathcal{P}$, we say y is between x, z if the cuts in x, z divide different elements of the coarse partition of $\mathcal{E}$ induced by y, and y is not equal to x, z.

Lemma 7.7. *This relation defines a pretree.*

Proof. Let

$$\mathcal{E} = x^{(1)} \sqcup \cdots \sqcup x^{(n_x)},$$
$$\mathcal{E} = y^{(1)} \sqcup \cdots \sqcup y^{(n_y)},$$
$$\mathcal{E} = z^{(1)} \sqcup \cdots \sqcup z^{(n_z)}$$

be the coarse partitions of the ends of Γ induced by x, y, z. First we check that the definition makes sense; i.e., given $x, y \in \mathcal{P}$, there are unique i, j with

$$\bigsqcup_{k\neq i} x^{(k)} \subseteq y^{(j)}.$$

If one of x, y is an equivalence class of totally isolated cuts, then Lemmas 7.4 and 7.6 yield this. Suppose both are #-classes Q, R. By Lemma 7.5, given $K \in R$ there is $Q^{(i)}$ such that

$$K^{(2)} \subseteq Q^{(i)}.$$

$Q^{(i)}$ is contained in a $\bar{Q}^{(i)}$, so

$$K^{(2)} \subseteq \bar{Q}^{(i)},$$

and furthermore this $\bar{Q}^{(i)}$ is independent of the cut K chosen. For each j, j', we can find $K \in R$ with $\bar{R}^{(j)}$, $\bar{R}^{(j')}$ in different $K^{(k)}$ since we are using the coarse partition, whence one of $\bar{R}^{(j)}$, $\bar{R}^{(j')}$ is contained in $\bar{Q}^{(i)}$. Hence all but one $\bar{R}^{(j)}$ is contained in $\bar{Q}^{(i)}$; i.e.,

$$\bigsqcup_{k\neq i} \bar{R}^{(k)} \subseteq \bar{Q}^{(i)}.$$

For part (1) of the definition of a pretree, note that if

$$\bigsqcup_{k\neq i_1} x^{(k)} \subseteq y^{(j_1)}, \quad \bigsqcup_{k\neq i_2} x^{(k)} \subseteq y^{(j_2)},$$

with $j_1 \neq j_2$, then since $y^{(j_1)}$, $y^{(j_2)}$ are disjoint, $n_x = 2 = n_y$ and x, y are equivalent cuts, and hence are equal as elements of $\mathcal{P}$. So xyx does not hold.

Part (2) is trivial. For part (3), after relabelling we have

$$\coprod_{k \neq 1} y^{(k)} \subseteq x^{(1)}, \quad \coprod_{k \neq 1} x^{(k)} \subseteq y^{(1)},$$

$$\coprod_{k \neq 2} y^{(k)} \subseteq z^{(1)}, \quad \coprod_{k \neq 1} z^{(k)} \subseteq y^{(2)}.$$

Then $y^{(1)} \cap y^{(2)} = \varnothing$ implies $x^{(i)} \cap z^{(j)} = \varnothing$ for $i, j \neq 1$. Hence

$$\coprod_{k \neq 1} x^{(k)} \subseteq z^{(1)},$$

so x, y divide the same $z^{(k)}$, and thus xzy does not hold.

For part (4), suppose that xzy so that

$$\coprod_{k \neq 1} z^{(k)} \subseteq x^{(1)}, \quad \coprod_{k \neq 2} z^{(k)} \subseteq y^{(1)};$$

i.e., x divides $z^{(1)}$, y divides $z^{(2)}$. If $w \neq z$ then w divides a unique $z^{(i)}$. If $i = 1$ then yzw. If not, then xzw. $\square$

We recall the vertex version of Menger's theorem (see, for instance, [Bondy and Murty 2008, Theorem 9.1]):

Menger's theorem. *Let Γ be a graph and a, b be vertices of Γ. Then the minimum size of a vertex cut separating a, b is equal to the maximum number of vertex-independent simple paths joining a, b.*

Lemma 7.8. *This pretree is discrete.*

Proof. Let K, L, M be mincuts with M between K, L. Elements of $\mathcal{P}$ are of course not cuts; take a representative cut of any equivalence class or an appropriate corner cut of a #-class. By Lemma 2.10, only finitely many cuts are not nested with both K, L, so we need only consider the case when M is nested with both. Let C_i, D_i, E_i denote components of $\Gamma - K$, $\Gamma - L$, $\Gamma - M$ respectively. We have that K is nested with M, so (after relabelling if necessary) $C_1 \subseteq E_1$, and similarly $D_1 \subseteq E_2$. We have $E_1 \neq E_2$ as M is between K, L. By the remarks following Definition 2.6, E_1 is contained in a component D_i, whence K, L are nested.

If we now form a new graph by collapsing both of K, L to a single vertex and apply Menger's theorem in this graph, we obtain n vertex-independent paths from K to L, where n is the cardinality of a mincut. In the case when K, L are not disjoint, some of these paths collapse into points. The cut M must intersect each of

Figure 16. A succulent.

these paths as it separates K, L, and $|M| = n$, so M is contained in the union of these paths. Then there are only finitely many choices for M.

If we took different choices for K, L the only additional choices for M would be equivalent in $\mathcal{P}$ to some already considered. So the pretree is discrete.   □

We now have a discrete pretree $\mathcal{P}$, which as discussed above gives us a tree encoding the mincuts of Γ and how they interact with the ends of Γ.

8. Succulents

We have now obtained a tree encoding the cuts of the graph, with #-classes collapsed down to points. We now seek to reintroduce the cyclic structure of these in order to obtain the final "cactus" theorem. We will not be able to use cactus graphs as such; these work well for encoding edge cuts, but cannot represent a vertex cut yielding several components. We will therefore use a slightly more general structure which, for the sake of a horticultural joke, we call succulents.

Definition 8.1. A *succulent* is a connected graph built up from cycles (including possibly 2-cycles, consisting of two vertices joined by a double edge) in the following manner. Two cycles may be joined together either at a single vertex or along a single edge. The construction is tree-like in the sense that if we have a "cycle of cycles" $C_1, \ldots, C_n$ with C_i attached to C_{i+1} (mod n) then all the C_i share a common edge/vertex. See Figure 16 for an example. The analogous property in a tree is that if we have a cycle of edges with each attached to the next one, they all meet at a common vertex. An *end vertex* of a succulent is one contained in only one cycle; a vertex of a succulent is an end vertex if it has at most two edges adjacent to it.

Figure 17. Diagrams illustrating how we connect rings around a star vertex if both corner cuts are B-cuts (left), if the corner cuts are nested and not equal (middle), and if they are equal (right).

We now construct a succulent encoding the mincuts of Γ. We already have the tree T from the previous section whose vertices are (equivalence classes of) totally isolated cuts and #-classes joined together via "star" vertices. There is at most one star for each corner cut of a #-class. If there is no corresponding star, then the components split off by this cut are not further subdivided by mincuts.

Before moving on further, we note that totally isolated cuts can be represented by a degenerate sort of ring, constructed by attaching triangles to a segment rather than to a cycle. So we can always talk about the anchors of a member of $\mathcal{P}$.

To form our succulent, we replace each member of $\mathcal{P}$ by its associated ring. We must now consider how we connect these; i.e., we need to consider the behaviour around each star vertex. Recall that if Q is a #-class attached to some star vertex, all members of $\mathcal{P}$ divide the same member of the (coarse) partition of the ends of Γ corresponding to Q, $\bar{Q}^{(1)}$ say, and that there is a corner cut of Q separating $\bar{Q}^{(1)}$ from the rest of the ends.

Suppose Q, R are #-classes adjacent to the same star vertex, so that there is no member of $\mathcal{P}$ between them. See Figure 17. Let K, L be the corresponding corner cuts and $\bar{Q}^{(1)}$, $\bar{R}^{(1)}$ the members of the coarse partition. If both K, L are B-cuts then each of $\bar{Q}^{(1)}$, $\bar{R}^{(1)}$ comprises only one $Q^{(i)}$, $R^{(i)}$ and there is only one member of the fine partitions divided by the other #-class. We join these classes by identifying the appropriate anchors. If there are no other elements of $\mathcal{P}$ joined to this star vertex then K, L are equivalent so we could further simplify things by removing the anchors and joining the cycles for Q, R together directly.

If one of K, L is not a B-cut then the two cuts are nested. Then either they are equal or all components except one of $\Gamma - L$ are contained in the same component of $\Gamma - K$ and vice versa. In the latter case, there is only one member of the fine partitions divided by the other #-class, so again we can represent this by identifying the appropriate anchors. If the corner cuts are equal, then we glue together the rings via the corner cuts. Of the anchors attached to each, one represents the other #-class and we simply delete this; the other anchors come in pairs, each representing the same set of ends but originating from different rings; we identify these together so we don't get redundancy. Then to produce our succulent, we first glue together

those rings sharing a corner cut, and then attach the other members of $\mathcal{P}$ adjacent to this star by identifying the appropriate anchors (for totally isolated cuts we simply note that the coarse and fine partitions coincide so there will be an obvious anchor to use and we have none of the issues above).

We must now show that this is a true succulent, that is, that we still have a tree-like structure. We inherit much of the tree-like nature from T; we only need check that no "cycles of cycles" form from the identifications made between rings all adjacent to the same star vertex. We will proceed by contradiction, supposing we have a shortest cycle of cycles $\mathcal{C}$ preventing our graph being a succulent. We can place limitations on which constituent cycles of the graph can be present in $\mathcal{C}$. First the cycles on which our rings are based do not appear. This is because any two cycles meeting one of these in the same $\bar{Q}^{(i)}$ intersect along an edge. So $\mathcal{C}$ consists of the triangles which contain anchors; these can be joined together either at an anchor or along the opposite edge. Because our cycle is shortest, we alternate between joins along edges and at anchors. Hence our cycle has at least four members. Let $T_1, \ldots, T_4, \ldots$ be the triangles in $\mathcal{C}$ with T_1, T_2 meeting at an anchor, T_2, T_3 at an edge and so on. By construction the points at the bases of the T_i represent cuts K_i of Γ partitioning the ends of Γ, and after suitable labelling we have

$$\coprod_{i \neq 1} K_1^{(1)} \subseteq K_2^{(1)} = K_3^{(1)},$$

$$\coprod_{i \neq 1} K_4^{(1)} \subseteq K_2^{(2)} = K_3^{(2)},$$

whence

$$\coprod_{i \neq 1} K_1^{(1)} \subseteq K_4^{(1)} \subseteq K_6^{(1)} \ldots.$$

But $\mathcal{C}$ is a cycle, so we eventually come back to the start, whence all the inequalities become equalities, K_1 becomes a B-cut and

$$K_1^{(2)} = K_2^{(1)} = K_4^{(1)} = \ldots.$$

We could have started at any other point, so the other K_i are also B-cuts and all of them are equivalent. Then $\mathcal{C}$ becomes trivial and we have indeed constructed a succulent.

We have now proved most of this:

Theorem 8.2. *Let Γ be a connected graph such that there are vertex end cuts of Γ with finite cardinality. There is a succulent $\mathcal{S}$ with the following properties*:

(1) *There is a subset A of vertices of $\mathcal{S}$ called the* anchors *of $\mathcal{S}$. If two anchors are adjacent, one of them is an end vertex of the graph. Every vertex of $\mathcal{S}$ not in A is adjacent to an anchor. We define an anchor cut of $\mathcal{S}$ to be a vertex cut*

containing no anchors which separates some anchors of $\mathcal{S}$. We say anchor cuts are equivalent if they partition A in the same way.

(2) *There is an onto map f from the ends of Γ to the union of the ends of $\mathcal{S}$ with the end vertices of $\mathcal{S}$ which are anchors.*

(3) *There is a bijective map g from equivalence classes of minimal end cuts of Γ to equivalence classes of minimal anchor cuts of $\mathcal{S}$ such that ends e_1, e_2 of Γ are separated by $[K]$ if and only if $f(e_1), f(e_2)$ are separated by $g([K])$.*

(4) *Any automorphism of Γ induces an automorphism of $\mathcal{S}$.*

Proof. We already have a succulent containing a representative of each mincut; i.e., we already have the map g. We now discuss how we modify the succulent to define the map f of the ends of Γ. Some issues arise because there may be ends of Γ which are distinguished from each other only by nonminimal cuts. If such ends exist, we will treat them as a single end for the present section; i.e., we will map them all to the same place using f. Let ϵ be an end of Γ. If there is a mincut K such that this end is the sole element of one of the sets $K^{(i)}$, then this mincut appears somewhere in the succulent either as a corner cut of a #-class or as a totally isolated cut and there is an end anchor of the succulent corresponding to this $K^{(i)}$; define $f(\epsilon)$ to be this anchor.

If not, there may be a sequence of $x_i \in \mathcal{P}$ with

$$x_1^{(1)} \supseteq x_2^{(1)} \supseteq \cdots \ni \epsilon.$$

This defines a ray in the tree T associated to $\mathcal{P}$, hence an end of that tree. There is a unique such end, since T is a tree so two ends can be separated using a single point, which we may take to be some $y \in \mathcal{P}$. But there is only one $y^{(i)}$ containing ϵ, so only one end will do. So we have an end of T, hence of the succulent, associated with ϵ; this is where we will map ϵ.

The remaining cases will correspond to ends which can only be split off by nonminimal cuts, which are not associated to some end of the tree T. To fit these into our succulent, we will essentially pretend that they can be split off by a mincut; we will add an element to $\mathcal{P}$ for each such end, inducing a partition

$$\mathcal{E} = \{\epsilon\} \cup (\mathcal{E} - \{\epsilon\}).$$

This member of $\mathcal{P}$ will not be between any two members of $\mathcal{P}$; and it does not disrupt the discreteness of $\mathcal{P}$ because an infinite set of betweenness in $\mathcal{P}$ would induce a descending sequence of partitions as above, so we would already have dealt with this end. So we have added an end vertex to the tree T. When constructing the succulent, the extra member of $\mathcal{P}$ will be modelled as two anchors joined by a double edge, one of which becomes attached to a relevant anchor in the succulent. The other anchor is an end anchor, which we define to be $f(\epsilon)$.

We now have the map f, which by construction interacts with g in the way stated; note that the extra anchors added in the third step above are never split off by an anchor cut of $\mathcal{S}$.

To see that f is onto, we note that any end anchors of $\mathcal{S}$ arise either as $f(\epsilon)$ in the third case above, or as part of a ring, where they correspond to some member of a partition of $\mathcal{E}$, whose members will be mapped there. Any ends of $\mathcal{S}$ arise from ends of T, hence from sequences of members of $\mathcal{P}$. From the vertices in the relevant cuts, we can construct a ray in Γ, giving an end that will be mapped to the end of $\mathcal{S}$.

Part (4) arises since an automorphism of Γ induces corresponding automorphisms of the cuts and ends of Γ, preserving crossings, nestings, equivalences; in short, all the information used to construct $\mathcal{S}$. □

We make some remarks about the theorem. In part (3) we must say equivalence classes of cuts of $\mathcal{S}$ because we may have equivalent distinct cuts of $\mathcal{S}$; these arise if there are quasiequivalent, nonequivalent half-cuts in a #-class , whence there will be some equivalent cuts contained in the relevant ring; but this is not really a concern.

If we wish to obtain a graph in which we do not have to exclude anchors from cuts, we can replace each anchor with a 3-vertex and treat these as ends, so that the anchor cuts in the theorem become bona fide mincuts of the resulting graph $\mathcal{S}'$.

If we collapse the extra end anchors we added in the proof above onto the adjacent anchors, then we obtain a variant theorem:

Theorem 8.3. *Let Γ be a connected graph such that there are vertex end cuts of Γ with finite cardinality. There is a succulent $\mathcal{S}$ with the following properties*:

(1) *There is a subset A of vertices of $\mathcal{S}$ called the* anchors *of $\mathcal{S}$. No two anchors are adjacent, and every vertex of $\mathcal{S}$ not in A is adjacent to an anchor. We define an anchor cut of $\mathcal{S}$ to be a vertex cut containing no anchors which separates some anchors of $\mathcal{S}$. We say anchor cuts are equivalent if they partition A in the same way.*

(2) *There is a map f from the ends of Γ to the union of the ends of $\mathcal{S}$ with the anchors of $\mathcal{S}$.*

(3) *There is a bijective map g from equivalence classes of minimal end cuts of Γ to equivalence classes of minimal anchor cuts of $\mathcal{S}$ such that ends e_1, e_2 of Γ are separated by $[K]$ if and only if $f(e_1)$, $f(e_2)$ are separated by $g([K])$.*

(4) *Any automorphism of Γ induces an automorphism of $\mathcal{S}$.*

Consider a finite graph Γ. We call a set J of vertices of a graph Γ *n-inseparable* if $|J| \geq n+1$ and for any set K of vertices with $|K| \leq n$, J is contained in a single component of $\Gamma - K$. Let κ be the smallest integer for which there are κ-inseparable sets J_1, J_2 and a vertex cut K with $|K| = \kappa$ and J_1, J_2 in different components of $\Gamma - K$. We can consider the maximal κ-inseparable sets of Γ as ends of the graph;

or attach a sequence of $(\kappa+1)$-vertices to each to turn them into a bona fide end. The inseparability conditions ensure that this does not affect the cuts of Γ of size κ. Then the size-κ vertex cuts separating two inseparable sets become minimal end cuts of our graph, so we can obtain a succulent theorem for them:

Theorem 8.4. *Let Γ be a finite connected graph such that there exists κ for which there are κ-inseparable sets J_1, J_2 and a vertex cut K with $|K| = \kappa$ and J_1, J_2 in different components of $\Gamma - K$, and take the minimal such κ. There is a succulent S with the following properties*:

(1) *There is a subset A of vertices of S called the* anchors *of S. No two anchors are adjacent, and every vertex of S not in A is adjacent to an anchor. We define an anchor cut of S to be a vertex cut containing no anchors which separates some anchors of S. We say anchor cuts are equivalent if they partition A in the same way.*

(2) *There is a map f from the κ-inseparable sets of Γ to the anchors of S.*

(3) *There is a bijective map g from equivalence classes of minimal cuts of Γ separating κ-inseparable sets to equivalence classes of minimal anchor cuts of S such that κ-inseparable sets J_1, J_2 of Γ are separated by $[K]$ if and only if $f(J_1)$, $f(J_2)$ are separated by $g([K])$.*

(4) *Any automorphism of Γ induces an automorphism of S.*

Tutte [1984] produced structure trees for the cases $\kappa = 1, 2$, which Dunwoody and Krön [2015] then extended to higher κ. These trees were based on "optimally nested" cuts in the language of [loc. cit.], which in this case means A-cuts. Roughly speaking, the trees consist of the totally isolated cuts and corner cuts of our succulents, together with "blocks" which are not decomposed by the cuts in question; these include the maximal inseparable sets, and also sets broken up by cuts which are not optimally nested; these sets correspond to the #-classes. The structure trees can then be obtained from our succulents by replacing each ring with a star with one central vertex and one vertex joined to it for each corner cut. So these earlier results also follow from our work.

9. Applications

First we note that our work yields a proof of Stallings' theorem, based on the Bass–Serre theory of groups acting on trees (see [Serre 1980]).

Stallings' theorem. *Let G be a finitely generated group acting transitively on a graph Γ with more than two ends. Then G can be expressed as an amalgam $G = A *_F B$ or an HNN extension $G = A*_F$, where F has a finite index subgroup which is the stabilizer of a vertex of Γ.*

Proof. From the pretree $\mathcal{P}$ we obtain a tree T on which G acts. The tree T is nontrivial; the action is transitive and Γ has more than two ends so there are infinitely many ends and many inequivalent cuts. The action is without inversion since T is bipartite, formed of star vertices and elements of $\mathcal{P}$. Then G is isomorphic to the fundamental group of a certain graph of groups; G is finitely generated so this graph is finite. The action is nontrivial as G acts transitively on Γ, so it follows that G splits over the stabilizer of an edge of T. An element fixing an edge of T fixes the adjacent element of $\mathcal{P}$, and hence fixes either a #-class or an equivalence class of totally isolated cuts. A #-class contains finitely many vertices; and the transitivity of the action implies that there can only be finitely many cuts in each equivalence class, since we can find two cuts between which every cut of the class lies, and then apply the methods of Lemma 7.8. The result follows.   □

Stallings' original theorem covers the two-ended case as well, but our tree is trivial here. The two-ended case can be covered by more elementary means, however.

We now discuss how earlier cactus theorems concerning edge cuts follow from ours. We turn a question about edge end cuts into a question about vertex end cuts as follows. First replace the graph Γ with its barycentric subdivision Γ^b. This is defined as follows:

$$V(\Gamma^b) = V(\Gamma) \cup E(\Gamma),$$
$$E(\Gamma^b) = \{(v, e) : v \in V(\Gamma), e \in E(\Gamma), v \text{ an endpoint of } e\}.$$

If the cardinality of a minimal edge end cut of Γ is n, we now "thicken up" each vertex of Γ^b that was a vertex of Γ by replacing it with an $(n+1)$-vertex (see Definition 6.1) to obtain a graph Γ^*. In this way, an edge cut of Γ separating some ends of Γ corresponds precisely with a vertex cut of Γ^* of the same cardinality. In Γ^*, because all the vertex cuts are essentially edge cuts, all of the minimal vertex cuts of Γ^* split the graph into precisely two pieces, each containing an end. So we do not need to remove slices from the graph, and all cuts are B-cuts. It follows that quasiequivalent half-cuts are equivalent, and each ring becomes simple enough to be replaced by a cycle, in which the anchors become the vertices and the other vertices become the edges. Our succulent from Theorem 8.3 can then be replaced with a cactus, so we have the cactus theorem for edge end cuts:

Theorem 9.1 [Evangelidou and Papasoglu 2014]. *Let Γ be a connected graph such that there are edge end cuts of Γ with finite cardinality. There is a cactus $\mathcal{C}$ with the following properties*:

(1) *There is a map f from the ends of Γ to the union of the ends of $\mathcal{C}$ with the vertices of $\mathcal{C}$.*

(2) *There is a bijective map g from equivalence classes of minimal end cuts of Γ*

to minimal edge cuts of $\mathcal{C}$ such that ends e_1, e_2 of Γ are separated by $[K]$ if and only if $f(e_1)$, $f(e_2)$ are separated by $g([K])$.

(3) *Any automorphism of Γ induces an automorphism of $\mathcal{C}$.*

To deal with the classical cactus theorem for edge cuts of finite graphs, we proceed as before to get the graph Γ^*. Then to each $(n+1)$-vertex we attach an infinite chain of $(n+1)$-vertices, so that a vertex in the original graph Γ becomes a de facto end of our new graph. "Equivalent cuts" in this graph correspond to the same cut of the original graph. Once again the succulent can be replaced with a cactus, so we have the cactus theorem of Dinic, Karzanov, and Lomonosov:

Theorem 9.2 [Dinic et al. 1976]. *Let Γ be a connected finite graph. There is a cactus $\mathcal{C}$ with the following properties*:

(1) *There is a map f from the vertices of Γ to the vertices of $\mathcal{C}$.*

(2) *There is a bijective map g from equivalence classes of minimal edge cuts of Γ to minimal edge cuts of $\mathcal{C}$ such that vertices v_1, v_2 of Γ are separated by $[K]$ if and only if $f(v_1)$, $f(v_2)$ are separated by $g([K])$.*

(3) *Any automorphism of Γ induces an automorphism of $\mathcal{C}$.*

Acknowledgements

References

[Bondy and Murty 2008] J. A. Bondy and U. S. R. Murty, *Graph theory*, Graduate Texts in Mathematics **244**, Springer, New York, 2008. MR 2009c:05001 Zbl 1134.05001

[Bowditch 1999] B. H. Bowditch, *Treelike structures arising from continua and convergence groups*, vol. 139, Mem. Amer. Math. Soc. **662**, Amer. Math. Soc., Providence, RI, 1999. MR 2000c:20061 Zbl 0961.20034

[Dinic et al. 1976] E. A. Dinic, A. V. Karzanov, and M. V. Lomonosov, "The structure of a system of minimal edge cuts of a graph", pp. 290–306 in Исследования по дискретной оптимизации (*Studies in discrete optimization*), edited by A. A. Fridman, Nauka, Moscow, 1976. In Russian. MR 58 #16395

[Droms et al. 1995] C. Droms, B. Servatius, and H. Servatius, "The structure of locally finite two-connected graphs", *Electron. J. Combin.* **2** (1995), Research Paper 17. MR 96f:05124 Zbl 0829.05041

[Dunwoody and Krön 2015] M. J. Dunwoody and B. Krön, "Vertex Cuts", *J. Graph Theory* **80**:2 (2015), 136–171.

[Evangelidou and Papasoglu 2014] A. Evangelidou and P. Papasoglu, "A cactus theorem for end cuts", *Internat. J. Algebra Comput.* **24**:1 (2014), 95–112. MR 3189668 Zbl 1286.05078

[Fleiner and Frank 2009] T. Fleiner and A. Frank, "A quick proof for the cactus representation of mincuts", technical report QP-2009-03, Egerváry Research Group, Budapest, 2009, available at http://www.cs.elte.hu/egres.

[Huntington 1935] E. V. Huntington, "Inter-relations among the four principal types of order", *Trans. Amer. Math. Soc.* **38**:1 (1935), 1–9. MR 1501800 Zbl 0012.14601

[Huntington and Rosinger 1932] E. V. Huntington and K. E. Rosinger, "Postulates for separation of point-pairs (reversible order on a closed line)", *Proc. Amer. Acad. Arts Sci.* **67** (1932), 61–145.

[Serre 1980] J.-P. Serre, *Trees*, Springer, Berlin, 1980. MR 82c:20083 Zbl 0548.20018

[Stallings 1968] J. R. Stallings, "On torsion-free groups with infinitely many ends", *Ann. of Math.* (2) **88** (1968), 312–334. MR 37 #4153 Zbl 0238.20036

[Tutte 1984] W. T. Tutte, *Graph theory*, Encyclopedia of Mathematics and its Applications **21**, Addison-Wesley, Reading, MA, 1984. MR 87c:05001 Zbl 0554.05001

GARETH R. WILKES
MATHEMATICAL INSTITUTE
UNIVERSITY OF OXFORD
ANDREW WILES BUILDING, RADCLIFFE OBSERVATORY QUARTER
WOODSTOCK ROAD
OXFORD
OX2 6GG
UNITED KINGDOM

gareth.wilkes@sjc.ox.ac.uk