CS 15-853: Algorithms in the Real World **October 09, 2002**

More on Planar Graph Separators

Lecturer: Bruce Maggs *Scribe: Hubert Chan*

Outline: In the lecture "Graph Separators II", we assumed the subroutine `CUTSHALLOW` takes a planar graph that has a spanning tree with not more than d levels, and returns a $\frac{2}{3}$-vertex separator of size at most $2d+1$. We show that such an algorithm exists by giving a constructive proof.

1 Introduction

Suppose we have a graph embedded into the plane such that the edges do not cross one another. Note that a simple cycle in the graph forms a vertex separator (though not necessarily a good one). It separates the vertices inside it from those outside it, since there cannot be any edges between vertices on different sides of the cycle. This is the approach[1] we use to find a suitable separator. We state the main result of the lecture.

Theorem 1.1 *Suppose a connected planar graph has a spanning tree whose depth is bounded by d. Then, the graph has a $\frac{2}{3}$-vertex separator[2] of size at most $2d+1$.*

Note:

1. We give a constructive proof of Theorem 1.1, in which an algorithm to find such a separator is outlined.
2. If the graph has n vertices and its diameter is bounded by $\sqrt{n}$, then a breadth-first search tree of the graph has depth bounded by $\sqrt{n}$ and hence by Theorem 1.1 it has a $\frac{2}{3}$-vertex separator of size at most $2\sqrt{n}+1$.
3. We assume that the graph is already embedded into a plane such that the edges do not cross one another.
4. Suppose the graph is not connected.
   (a) If it has a component of size between $\frac{n}{3}$ and $\frac{2n}{3}$, then we can obtain a $\frac{2}{3}$-separation by grouping all other components together.

[1] This approach is also outlined in Satish Rao's lecture notes from the course CS270 Combinatorial Algorithms & Data Structures that was taught in UC Berkeley.

[2] Recall a $\frac{2}{3}$-separation means that each of the separated components does not contain more than $\frac{2n}{3}$ vertices.

(b) If all the components have less than $\frac{n}{3}$ vertices, we form a set of vertices by adding the components to it, one at a time. Then, there exists a stage when the size of the set is between $\frac{n}{3}$ and $\frac{2n}{3}$. By the previous step, we are done.

(c) Suppose there exists a component of size greater than $\frac{2n}{3}$. We apply Theorem 1.1 to that component. Then, the larger sub-component has size at least $\frac{1}{2} \cdot \frac{2n}{3} = \frac{n}{3}$ and at most $\frac{2n}{3}$.

We conclude this section by giving an overview of our proof. We first fix some spanning tree T of the graph that has depth bounded by d. The graph is then triangulated to produce an augmented graph G. Then, we consider a subgraph D of the dual of G such that the arcs of D are exactly those that do not cross the edges in T. Since G is triangulated, the nodes in its dual have degrees bounded by three, and so do those in D. We observe that D is a tree and every arc in D corresponds to a non-tree edge in G, which induces a unique cycle when added to T. One such cycle would be a suitable vertex separator. In order to compute the number of vertices enclosed by such cycles, we assign weights to the nodes in D such that the sum of the weights of nodes in a subtree equals the number of vertices enclosed by the corresponding cycle. Finally, we show that a suitable cycle can be found conveniently using a $\frac{2}{3}$-arc separator for the labeled tree in D, which has node degree bounded by three.

2 Proof of Theorem 1.1

Suppose V is the set of vertices of the given graph and let T be the set of edges of a spanning tree with depth bounded by d. We triangulate[3] the graph by adding edges and let U be the set of edges that are added or not included in T. (Hence, $T \cap U = \phi$.) Let the augmented graph be $G := (V, E)$, where $E = T \cup U$.[4]

Let N be the set of (triangular) faces of G, or equivalently,

$$N := \{X \mid X \subseteq V,\ X \text{ forms a face (triangle) in } G\}.$$

Let H be the set of pairs of faces that share an edge in E, or equivalently,

$$H := \{\{X, Y\} \mid X, Y \in N,\ X \cap Y \in E\}$$

The *dual* of G is defined to be the graph (N, H), i.e., a graph in which each face is a node and there is an arc between any pair of faces that share an edge in E.[5]

In this proof, we consider a subgraph D of the dual such that the arcs do not cross the edges in T. Let A be the set of pairs of faces that share an edge in U, or equivalently, $A := \{\{X, Y\} \mid X, Y \in N,\ X \cap Y \in U\}$. Then, $D := (N, A)$. Note that there is a natural 1-1 correspondence between the sets U and A.

[3] We can assume without loss of generality that $|V| \geq 3$.

[4] By Euler's Theorem for connected triangulated planar graphs, we have $|T \cup U| = 3|V| - 6$.

[5] To avoid confusion with graph G, we call an element in N a *node* (which is also called a face) and that in H an *arc*.

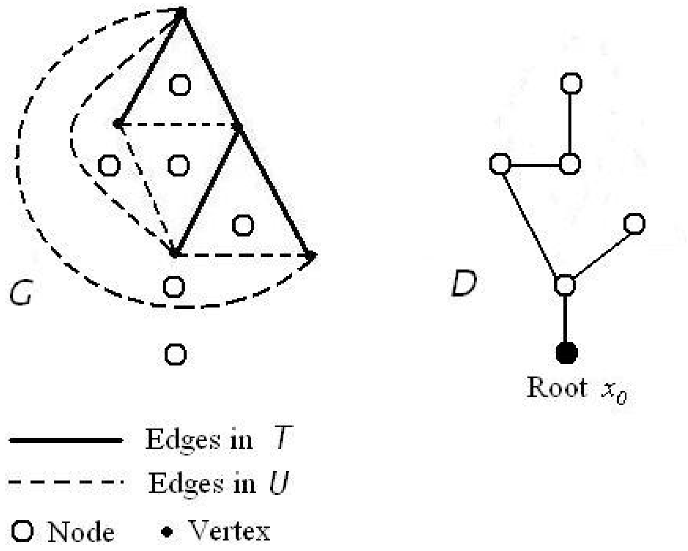

Figure 1: An Example of Graph G and Its Dual D

Lemma 2.1 *The set A of arcs forms a spanning tree for D.*[6]

Proof: Suppose there exists $x, y \in N$ that are not connected in D. This means we cannot traverse from face x to face y without crossing the edges of T. However, this implies that T contains a cycle, which is a contradiction. Hence, A spans N.

Suppose D contains a cycle. Then, there exist vertices in V on both sides of the cycle, since each arc in the cycle corresponds to an edge in U that it crosses. However, this implies that the vertices on different sides of the cycle cannot be connected by the edges in T, which is a contradiction, since T spans V.

■

2.1 Some Properties of G and D

Since D does not contain any cycle, there exists a node with degree one. We arbitrarily pick such a node x_0 and call it the *root* of tree A. We use terms like "child", "parent" and "leaf" in their usual sense associated with a rooted tree. Given a cycle in G, we define the side containing x_0 to be *outside*.

Since T is a spanning tree for G, adding an edge e in U to T induces a unique cycle C in G. Since the depth of T is bounded by d, the length of the cycle is at most $2d + 1$. Recall that edge e corresponds to some arc f in A crossing it. Suppose x is the child end of arc f in tree A, i.e., x is connected to its parent via arc f. We say C is the cycle *induced by x*. Note that every node in N, apart from the root x_0, induces a cycle in G (Figure 2). One of the cycles induced by the nodes will be our vertex separator.

[6]Using Euler's Theorem, we can show $|N| = 2|V| - 4$ and $|A| = 2|V| - 5$.

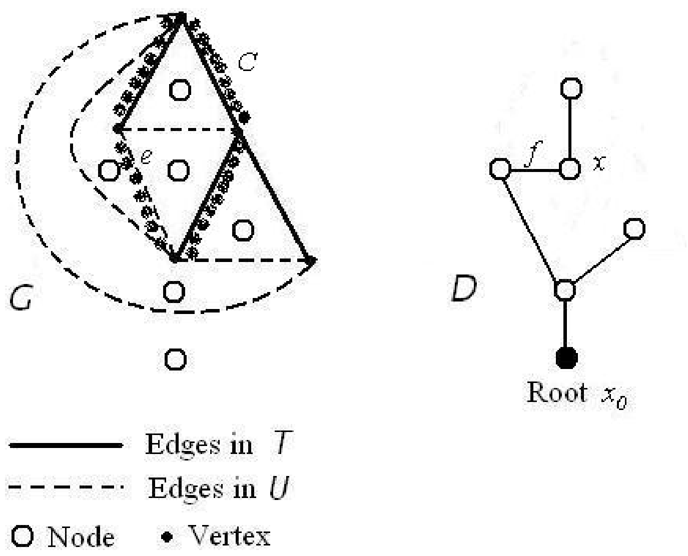

Figure 2: Correspondence between Node x, Arc f, Edge e and Cycle C

Given $W \subset V$, we say W is *contained* in a cycle C if all the vertices in W are either inside or on C. (Recall that *inside* is defined with respect to x_0, which is considered to be *outside*.) Note that if the vertices on a cycle C_1 are contained in a cycle C_2, then any vertex inside C_1 is also inside C_2. Similarly, we say a cycle C *encloses* W (which can be a set of nodes or a set of vertices) if all the elements in W are inside C. We have the following properties of cycles induced by the nodes in N.

Lemma 2.2 *Suppose C_v is the cycle induced by node v. Then, all the nodes in the subtree rooted at v are inside C_v.*

Proof: Since v induces a cycle, it has a parent, say node u. Since arc $\{u, v\}$ corresponds to some edge f in U that induces C_v, u and v are on different sides of C_v. Note that f is the only cycle edge that is also in U. Suppose w is a descendant of v. Then, there exists a path in A that connects them together. This implies that we can traverse from face v to face w without crossing edge f or any edge in T. Hence, w is on the same side of C_v as v. Similarly, the root x_0 is on the same side of C_v as u. Hence, all the nodes in the subtree rooted at v are inside C_v. ∎

Lemma 2.3 *The cycle induced by a node contains the cycle induced by any one of its children. Moreover, the cycles induced by sibling nodes do not enclose any common vertex.*

Proof: Suppose $u = \{a, b, c\} \in N$ is neither the root nor a leaf. Note that since G is triangulated, u can have either one or two children.

1. Suppose u has only one child, say v. (Figure 3)

   Suppose $\{b, c\}$ is the edge crossed by arc $\{u, v\}$. Note that the degree of u in D is two. Hence, without loss of generality, we have $\{a, b\} \in T$ and $\{b, c\}$, $\{a, c\} \in U$. Let C_u and C_v by the cycles induced by u and v respectively.

   (a) Suppose a does not lie on the cycle C_v.
   Since $a, b \in T$, the cycle C_u can be traced by starting at a, first traversing $\{a, b\}$, then all the vertices on C_v, until reaching c, and finally traversing the edge $\{c, a\}$ to return to a. Hence, C_u contains C_v, and the two cycles enclose the same vertices in V.

   (b) Suppose a lies on the cycle C_v.
   The cycle C_v can be traced by starting at b, traversing $\{b, a\}$, and then along path p in the cycle to reach c and then back to b. Note that C_v contains exactly one edge in U and hence does not contain $\{c, a\}$, and path p consists of edges in T.
   Consider the cycle C_u induced by the edge $\{a, c\}$. Then, we can traverse C_u by starting at a, along p to reach c, and then traversing $\{c, a\}$ to complete the cycle. Since u is the parent of v, the root x_0 and v must be on different sides of C_u by Lemma 2.2. Hence, C_u contains C_v, and C_u encloses one extra vertex b.

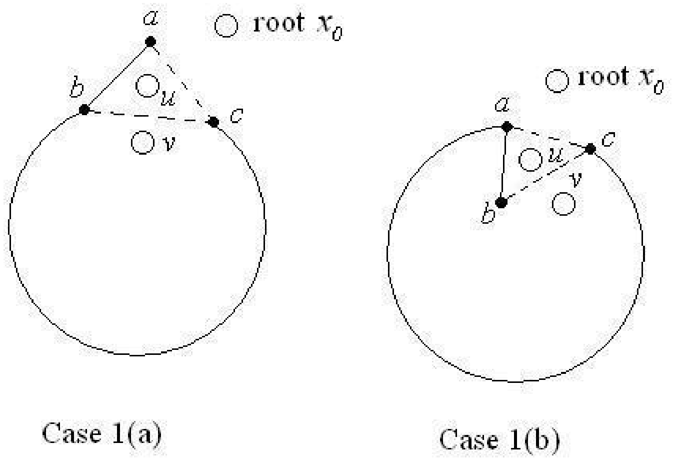

Figure 3: Node u has only one child v.

2. Suppose u has two children v and w (Figure 4), and let C_u, C_v and C_w be the cycles induced by the corresponding nodes.

Note that all three edges of the face u are in U. Without loss of generality, suppose arc $\{u, w\}$ crosses edge $\{a, b\}$, and $\{u, v\}$ crosses $\{b, c\}$.

Again by Lemma 2.2, neither v nor w can be on the same side of C_u as the root x_0, as they are the children of u. Hence, b is contained in C_u. Since T is a spanning tree of G, there exists a unique path p in T that connects b to a unique point g on the cycle C_u (g could be a or c). Note that the vertices on path p are contained in C_u.

It follows C_w can be traced by starting at a, along C_u to g, and then along p to b, and finally traversing $\{b, a\}$ back to a. Hence, C_w is contained in C_u. Similarly, C_v is contained in C_u. Recall G is planar and so p partitions the vertices enclosed in C_u. Hence, there is no vertex that is enclosed by both C_v and C_w.

Furthermore, in addition to the vertices enclosed by C_v or C_w, C_u also contains the vertices on the path p, and hence the number of extra vertices it encloses equals the length of p. Since the depth of T is bounded by d, the length of p is bounded by $2d$.

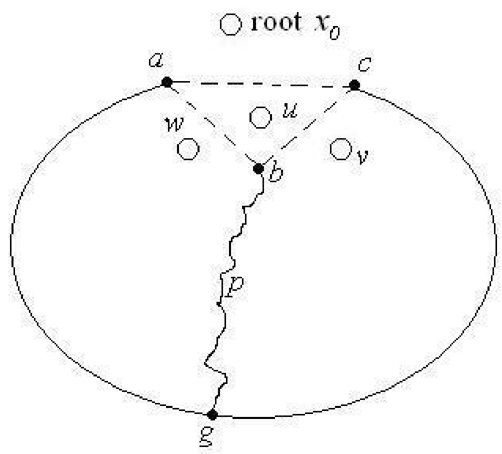

Figure 4: Node u has two children v and w.

■

2.2 Assigning Weights to the Nodes of Tree A

In order to evaluate how balanced the separation obtained by a cycle is, we need a way to calculate the number of vertices inside the cycle induced by a node in N. Because of Lemma 2.3, we have the following result.

Lemma 2.4 *We can assign a non-negative integer weight $W(u)$ to each node*[7] *u such that for all $u \in N \backslash \{x_0\}$, the number of vertices inside the cycle induced by u equals the sum of the weights of the nodes in the subtree rooted at u. We set the weight of x_0 to 0 by default. Moreover, the weight of a node is bounded by $2d$.*

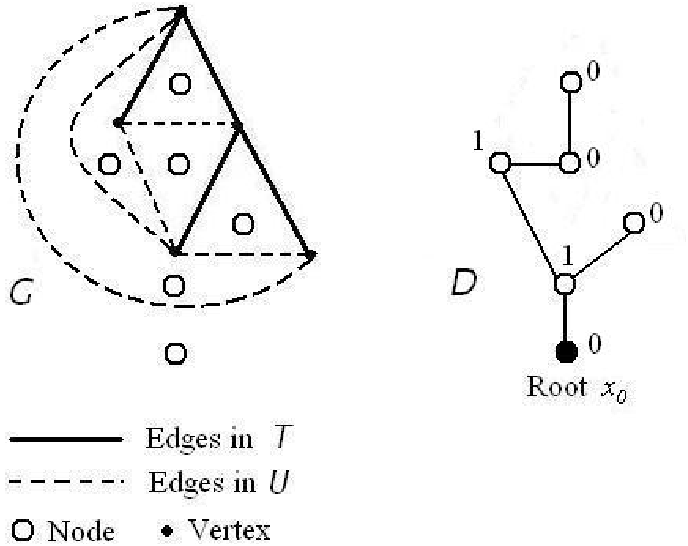

Figure 5: Assigning Weights to the Nodes of Tree A

Proof:

We describe an algorithm to assign weights to the nodes such that the lemma holds.

Suppose u is a leaf node in A. It follows that the face u has two edges in T and one edge in U. Hence, the cycle induced by u is a triangle in G and does not enclose any vertex. Every leaf node has weight 0.

Suppose u is neither the root nor a leaf. We analyze the cases in the same manner as in the proof of Lemma 2.3.

1. Node u has exactly one child v.

   (a) In this case, the cycle C_u does not enclose any vertex other than those already enclosed by C_v. Hence, the weight of u is 0.

   (b) In this case, the C_u encloses one more vertex in addition to those enclosed in C_v. Hence, the weight of u is 1.

[7]Note that instead of assigning a weight to a node (other than the root), it is equivalent to assign the weight to the arc connecting it to its parent.

2. Node u has two children v and w.

   The cycle C_u encloses the vertices inside C_v or C_w, and the vertices on path p that are inside C_u. Note that it is important that C_v and C_w do not enclose any common vertex. Hence, the weight of u is the length of p, which is bounded by $2d$.

■

Lemma 2.5 *After assigning the weights in the manner described in Lemma 2.4, the sum of the weights of all the nodes in N is $|V| - 3$.*

Proof: Since $W(x_0) = 0$, the sum of all the weights of the nodes equals the number of vertices inside the cycle induced by the only child of the root x_0. Since the root has degree one by choice, the cycle induced by the root's only child is a triangle corresponding to the face x_0. Hence, all the vertices apart from the three forming the triangle are *inside* the triangle. ■

2.3 Finding a Suitable Cycle Separator

Suppose the nodes in N are assigned weights in the manner described in Lemma 2.4. Removing an arc f in A separates the graph D into two connected components. The weight of a component is the sum of the weights of all the nodes in that component. We assign a direction to arc f such that it points to the side with heavier weight. Ties are resolved arbitrarily.

We say a node is *terminal* if it has no outgoing arcs. Note that the existence of a terminal node is guaranteed. Otherwise, we can visit the nodes in D along the arcs in the assigned directions without ever stopping. This contradicts the fact that D has no cycles. However, it is not important whether there exists a unique terminal node.

As a small digression, we consider a general result[8] whose proof exhibits interesting techniques.

Theorem 2.6 *Suppose each node in N is assigned a non-negative weight and K is a set of arcs that forms a tree spanning N. Let the sum of the weights of the nodes in N be S. Then, there exists a node whose removal together its incident arcs separates the graph (N, K) into two components, each of which has weight bounded by $\frac{2S}{3}$.*

Proof: Suppose we direct the arcs in K according to the above scheme and suppose u is a corresponding terminal node. We show that u is the node to be removed.

Note that the removal of u together with its incident arcs separates the graph into connected components, the number of which equals the degree of u. By the definition of a terminal node, the arcs joining u to the connected components all direct to u. Hence, each of the connected components has weight bounded by $\frac{S}{2}$.

[8] This is the well-known $\frac{2}{3}$-node separator for trees with weighted nodes.

1. Suppose there exists a connected component with weight between $\frac{S}{3}$ and $\frac{S}{2}$. Then, the sum of weights of all other connected components is bounded by $\frac{2S}{3}$. Hence, we have a separation satisfying the requirements.

2. Suppose each of the connected components has weight less than $\frac{S}{3}$.

   (a) If the sum of weights of the connected components is bounded by $\frac{2S}{3}$, then any way of grouping the connected components gives a suitable separation.

   (b) If the sum of weights of the connected components is larger than $\frac{2S}{3}$, we form a component by adding the connected components to it, one at a time. Then, there exists some stage at which the size of the component is between $\frac{S}{3}$ and $\frac{2S}{3}$. All the other remaining connected components are grouped together in another component, giving a satisfying separation.

■

Note that we cannot use Theorem 2.6 directly to find a suitable cycle separator. What we really need is not a single node separator, but a single arc separator. In general, a graph with unbounded degree does not have a single arc separator that gives a $\frac{2}{3}$-weight separation. However, since the graph G is triangulated, the degree of a node in N is at most three, and hence the argument can be modified.

Theorem 2.7 *Suppose the set K of arcs forms a spanning tree for N such that each arc has a non-negative weight. Let the sum of the weights of the arcs in K be S. If the degree of (N, K) is bounded by three, there exists an arc whose removal separates the graph (N, K) into two components, each of which has weight bounded by $\frac{2S}{3}$.*[9]

Proof:

Again, we pick a node with degree one to be the root x_0. Then, we have a rooted tree. We relabel the tree by pushing the weight of each arc to its child end, and assign weight zero to x_0. The arcs are directed as before and suppose u is a terminal node.

Suppose u has degree one. This means u has weight at least $\frac{S}{2}$ (and therefore cannot be the zero-weighted root x_0). Hence, the arc joining u to its parent gives a trivial separation, in which there is only one component with weight bounded by $\frac{S}{2}$.

Suppose u has degree two. As in the proof of Theorem 2.6, the removal of u separates the graph into two components, each of which has weight bounded by $\frac{S}{2}$. Consider the arc-weighted graph again. It follows the removal of the arc connecting u to its parent gives the same weight separation.

Suppose u has degree three. Note that the removal of u again separates the graph into three components, each of which has weight bounded by $\frac{S}{2}$. (1) Suppose each of the two child components has weight bounded by $\frac{S}{3}$. We can group the two child components together to

[9]The result can be generalized such that if the degree of each node in the tree is bounded by δ, there exists a single arc-separator that gives a $(1 - \frac{1}{\delta})$-weight separation.

form a $\frac{2}{3}$-node weight separation. Consider the arc-weighted graph. It follows the removal of the arc connecting u to its parent gives the same weight separation. (2) Suppose one of the child components has weight larger than $\frac{S}{3}$, say the subtree rooted at a child v of u. Consider the arc-weighted graph. Suppose we remove the arc $\{u, v\}$. Then, the connected component containing u would have weight less than $\frac{2S}{3}$, and the connected component containing v would still have weight bounded by $\frac{S}{2}$. Hence, we are done.

■

Nevertheless, we need a subtler argument to show that the above arc separator for the labeled tree indeed gives a suitable cycle separator for the original graph. Although the removal of a terminal node gives a $\frac{2}{3}$-weight separation, we have not proved that the number of vertices outside an induced cycle corresponds to the weights of the nodes in the tree in which the corresponding subtree is removed. Note that the weight of the arc (or its child-end node) removed could contribute to the number of vertices inside or outside the cycle, giving a possibility of upsetting the $(\frac{1}{3}, \frac{2}{3})$ balance. However, we can show that either (i) the weight of a terminal node is small (0 or 1), in which case it does not matter where it is included, or (ii) in case the weight of the terminal node is large, all its weight contributes to the number of vertices on the corresponding cycle separator. In either case, we show that the corresponding cycle is a $\frac{2}{3}$-vertex separator.

We prove the following lemma in a way similar to the proof of Theorem 2.7. This completes the proof of Theorem 1.1:

Lemma 2.8 *Consider (N, A) in which the arcs are directed according to the described way. Suppose u is a terminal node. Then, the cycle induced by u or one of its children gives a $\frac{2}{3}$-vertex separator of G that has size bounded by $2d + 1$.*

Proof: Let $|V| = n$. Note that the cycle induced by a node has length bounded by $2d + 1$ and so it suffices to prove that the separation obtained is balanced in the defined manner.

Since $W(x_0) = 0$, the root x_0 cannot be a terminal node. For the same reason, u cannot be a leaf. Let S be the sum of the weights of the nodes in N. (By Lemma 2.5, $S = n - 3$.) Again, we analyze the cases following the manner of Lemma 2.3:

1. Node u has exactly one child v.

   Let P be the subtree rooted at v. Since an arc directs from v to u, we have $W(P) \leq \frac{S}{2}$. Similarly, the fact that an arc is directed from u's parent to u implies $W(P) + W(u) \geq \frac{S}{2}$. Since $W(u)$ equals 0 or 1 in this case, we have

   $$\frac{n}{2} - \frac{5}{2} \leq \frac{S}{2} - W(u) \leq W(P) \leq \frac{S}{2} = \frac{n}{2} - \frac{3}{2}.$$

   The cycle induced by v gives a $\frac{1}{2}$-separation.

2. Node u has two children v and w.

   Let P and Q be the subtrees rooted at v and w respectively. Since u is a terminal node, we have $W(P) \leq \frac{S}{2}$, $W(Q) \leq \frac{S}{2}$ and $W(P) + W(Q) + W(u) \geq \frac{S}{2}$.

(a) Suppose $W(P) + W(Q) + W(u) \leq \frac{2n}{3}$.
Then, $\frac{n}{2} - \frac{3}{2} \leq W(P) + W(Q) + W(u) \leq \frac{2n}{3}$. Hence, the cycle induced by u gives a $\frac{2}{3}$-separation.

(b) Suppose $W(P) + W(Q) + W(u) > \frac{2n}{3}$.
Since $W(P) + W(Q) > \frac{2n}{3} - W(u)$, without loss of generality, we assume $W(P) \geq \frac{n}{3} - \frac{W(u)}{2}$. The number of vertices inside the cycle C_v is $W(P) \leq \frac{n}{2} - \frac{3}{2}$. Recall that in the proof of Lemma 2.3, the path p lies on the cycle C_v and hence we know C_v has length at least that of path p. Moreover, the length of p equals $W(u)$.

Hence, there are at least $W(u)$ vertices on C_v. The number of vertices outside C_v is hence bounded by $n - W(u) - W(P) \leq \frac{2n}{3} - \frac{W(u)}{2}$. The cycle C_v gives a $\frac{2}{3}$-separation.

■