Separators in Graphs with Negative and Multiple Vertex Weights

Hristo N. Djidjev* John R. Gilbert†

April 1994

Abstract

A separator theorem for a class of graphs asserts that every graph in the class can be divided approximately in half by removing a set of vertices of specified size. Nontrivial separator theorems hold for several classes of graphs, including graphs of bounded genus and chordal graphs.

We show that any separator theorem implies various weighted separator theorems. In particular, we show that if the vertices of the graph have real-valued weights, which may be positive or negative, then the graph can be divided exactly in half according to weight. If k unrelated sets of weights are given, the graph can be divided simultaneously by all k sets of weights.

These results considerably strengthen earlier results of Gilbert, Lipton, and Tarjan: (1) for $k = 1$ with the weights restricted to be nonnegative, and (2) for $k > 1$, nonnegative weights, and simultaneous division within a factor of $(1 + \epsilon)$ of exactly in half.

Keywords: graph separator, cutset, divide-and-conquer algorithm, separator theorem.

AMS(MOS) subject classifications: 05C10, 05C75, 68R10.

Computing Reviews descriptors: F.2.2, G.2.2.

*Department of Computer Science, Rice University, P.O. Box 1892, Houston, Texas 77025.

†Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, California 94304.

1 Introduction

A separator in a graph is a small set of vertices or edges whose removal divides the graph into pieces of approximately equal size. Graph separators have applications in partitioning for parallel computation [8], divide-and-conquer algorithms for sparse Gaussian elimination [7, 15], VLSI layout [13], geometric algorithms [17, 22], and computational complexity [4, 14].

A class of graphs is said to satisfy an $f(n)$-*separator theorem* (for vertices) if there are constants $\alpha < 1$ and $\beta > 0$ such that every n-vertex graph in the class has a set, called a *separator*, of at most $\beta f(n)$ vertices whose removal leaves no connected component with more than αn vertices. General graphs, and even sparse graphs, do not satisfy a nontrivial separator theorem [16]; but several useful classes of graphs do. Lipton and Tarjan [16] proved that planar graphs satisfy an $n^{1/2}$-separator theorem. Separator theorems are also known for trees [11], outerplanar graphs [13], chordal graphs [6], grid graphs [16, 23], hypercubes [9], several interconnection graphs for parallel computation [10, 12, 13], graphs of bounded genus [2, 5], graphs that exclude a given minor [1], and some geometrically defined classes that include many three-dimensional finite element graphs [19, 20, 21, 22].

In some applications it is useful to associate *weights* with the vertices of the graph. For example, weights can model different amounts of work in parts of a computation to be divided among parallel processors, or different sizes of circuit elements to be laid out. One may wish to divide a graph evenly by more than one kind of weight. A circuit might be partitioned both by layout area and power requirements, for example. A parallel task might be partitioned by its requirements for work, space, and communication. In this paper we show that an ordinary separator theorem for a class of graphs implies several kinds of weighted separator theorems, provided that the class includes all subgraphs of its members.

In the next section we review earlier work on weighted separator theorems. We also prove a "warmup" theorem saying that, if a graph's vertex weights add up to zero, then it has a separator of size $O(n^\lambda)$ that splits it into two pieces of nearly equal size and nearly zero weight. This significantly extends a result of Lipton and Tarjan [16], in which only nonnegative weights are allowed. The main technique in this proof is used in the more complicated and general results in later sections. In Section 3 we prove a technical lemma that is used in Section 4 to prove the most general result of the paper, which says that any number of weights, positive or negative, can be divided evenly in half by the same separator. The algorithm to find this separator has the same asymptotic running time, up to a constant factor, as the time required for the nonweighted separator algorithm. Section 5 contains concluding remarks.

This is a heavily revised version of an earlier technical report [3], in which we proved Theorem 3 (one mixed-sign weight) and the case $k = 2$ of Corollary 6 (two nonnegative weights).

Until further notice we shall take S to be a class of graphs that contains all subgraphs of its members, and that satisfies an n^λ-separator theorem for some fixed $\alpha < 1$, $\beta > 0$, and $0 < \lambda < 1$. When the graph G is clear from context, we shall often not distinguish between a subset of the vertices of G and the subgraph of G induced by those vertices. We write the size of a set A as $|A|$. If the weight of a vertex v is $\mathrm{wt}(v)$, we write the sum of the weights of a set A of vertices as $\mathrm{wt}(A)$. We define $\mathrm{sign}(A)$ to be -1, 0, or 1 according to the sign of $\mathrm{wt}(A)$.

2 Preliminaries

An n^λ-separator theorem requires that the separator divide the graph into connected components of a certain maximum size. It does not require that those components can be collected into exactly two groups of a certain maximum size. For example, three components of size $n/3$ cannot be grouped into two sets of size at most $n/2$.

Lipton and Tarjan [16] showed that an $n^{1/2}$-separator theorem for all subgraphs of G implies the apparently stronger fact that G has a separator of $O(n^{1/2})$ vertices that divides it into exactly two pieces, each with at most $n/2$ vertices. The pieces themselves might not be connected. We describe a generalized version of their algorithm that we will use to prove several other results.

We can think of the algorithm as accumulating the two pieces A and B and the separator C in buckets that are initially empty. At each iteration, the remaining portion of the graph is separated. One of the fragments is added to one of buckets A and B, the separator is added to bucket C, and the other fragment remains to be separated again. The iteration ends when the remaining fragment is empty; then C is a separator that divides A from B.

Figure 1 gives this algorithm in detail. It has blanks to be filled in with the separator algorithm to be used on the fragments, the method of choosing which fragment to put in which bucket, and some extra initialization.

```
procedure generic-separate(graph G);
    D ← all vertices of G;
    A, B, C ← ∅;
    (1): other initialization;
    while D ≠ ∅ do
        {invariant: A, B, C, D partition the vertices of G;
            no edge joins any two of A, B, D;
            (2): other invariants}
        (3): (X, Y, Z) ←separate(D);
        C ← C ∪ Z;
        (4): put at least one of X or Y into A or B;
        D ← the other one of X or Y (if any)
    end while
end procedure
```

Figure 1: A generic separator algorithm.

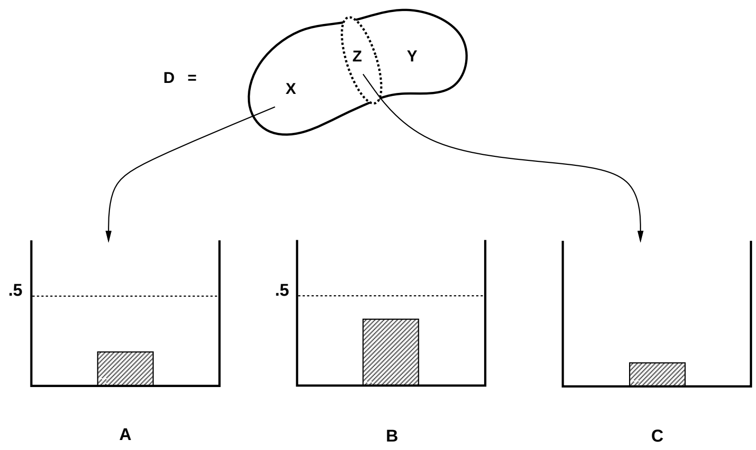

Figure 2: One iteration of the 50-50 weighted separator algorithm.

Lemma 1 *Let S be a class of graphs closed under subgraph and satisfying an n^λ-separator theorem. Suppose algorithm* separate(G) *runs in time $t(n) = \Omega(n)$ on every n-vertex graph in S, returning a partition (X, Y, Z) of the vertices, with $|X| \leq \alpha n$, $|Y| \leq \alpha n$, $|Z| \leq \beta n^\lambda$, and no edge joining X and Y. Then algorithm* generic-separate *runs in time $O(t(n))$ on every n-vertex graph G in S, and returns sets A, B, and C such that*

- *A, B, and C partition the vertices of G.*
- *No edge joins A and B.*
- $|C| = O(n^\lambda)$,

provided that Step 1 does not destroy the invariant and Steps 1, 3, and 4 all run in time $O(t(|D|))$.

Proof: If Step 4 does what it says, the algorithm will terminate because D gets strictly smaller at each iteration. In fact $|D|$ decreases geometrically, so the total running time will be $O(t(n))$ and the total size of C will be $O(n^\lambda)$. If the body of the loop preserves the "other invariants," then they too are true of the final partition. □

All the separator theorems in this paper will be proved by filling in the details of *generic-separate* in various ways. To illustrate, we first show that an n^λ-separator theorem for sizes implies an n^λ-separator theorem for nonnegative weights. Figure 2 is a diagram of the proof of this theorem.

Theorem 2 (50-50 weighted separator theorem [9, 16]) *Let G be an n-vertex graph whose subgraphs all satisfy an n^λ-separator theorem with parameters $\alpha < 1$ and $\beta > 0$. Suppose the vertices of G have nonnegative real weights whose sum is s. Then the vertices of G can be partitioned into sets A, B, and C such that C separates A and B, $|C| \le \delta n^\lambda$, and neither A nor B has weight more than $s/2$. Here δ is a constant that depends only on α and β. If the original n^λ-separator separator can be found in time $t(n)$, the 50-50 weighted separator can be found in time $O(t(n))$.*

Proof: First normalize the weights so that their sum is 1. Fill in the details of algorithm *generic-separate* as follows.

(1) No additional initialization.

(2) Additional invariant: $\mathrm{wt}(A) \le \mathrm{wt}(B) \le \frac{1}{2}$.

(3) At each iteration, separate the remaining vertices with the original n^λ-separator algorithm.

(4) To put one of X and Y into one of A and B:
exchange X and Y if necessary to make $\mathrm{wt}(X) \le \mathrm{wt}(Y)$;
$A \leftarrow A \cup X$;
exchange A and B if necessary to make $\mathrm{wt}(A) \le \mathrm{wt}(B)$.

Step 4 preserves the new invariant $\mathrm{wt}(A) \le \mathrm{wt}(B) \le \frac{1}{2}$ because $0 \le \mathrm{wt}(A) + \mathrm{wt}(X) \le \mathrm{wt}(B) + \mathrm{wt}(Y)$ and the sum of all the weights is one. At the end, the invariant says that neither A nor B has weight more than $1/2$ (or $s/2$, unnormalized); Lemma 1 gives the rest of the conclusion. $\Box$

We finish this section by proving a new theorem in which both positive and negative vertex weights are allowed, and the total weight is zero. Later we will prove a much stronger result (as Corollary 7), but this proof illustrates the technique used in the key step, Lemma 4.

What does it mean to separate a graph whose total weight is zero? A good separator should divide the graph into two parts both weighing approximately zero. But this is trivial: The empty separator suffices, where one part contains all the vertices and the other part is empty. Therefore we will also require that each part also contain at most some fixed fraction of the vertices. For the present, we will take "approximately zero" to mean "bounded by the magnitude of the weight of the separator."

Theorem 3 (Mixed-weight separator theorem) *Let G be an n-vertex graph whose subgraphs all satisfy an n^λ-separator theorem with parameters $\alpha < 1$ and $\beta > 0$. Suppose the vertices of G have real weights* $\mathrm{wt}(v)$ *whose sum is zero. Then the vertices of G can be partitioned into sets A, B, and C such that*

- $|\mathrm{wt}(A)| \le |\mathrm{wt}(C)|$ *and* $|\mathrm{wt}(B)| \le |\mathrm{wt}(C)|$,
- *No edge joins A and B.*
- $|A| \le \gamma n$, $|B| \le \gamma n$, *and* $|C| \le \delta n^\lambda$.

Here $\gamma < 1$ and δ are constants that depend only on α and β. If the original n^λ-separator separator can be found in time $t(n)$, the mixed-weight separator can be found in time $O(t(n))$.

Proof: Fill in the details of algorithm *generic-separate* as follows.

(1) Intially, call subroutine *separate* once on G and then once again on the larger of the resulting pieces, thus partitioning the vertices of G into a separator Z with at most $2\beta n^\lambda$ vertices and three pieces X_1, X_2, and X_3 such that no edge joins X_i to X_j if $i \ne j$, and such that $|X_i \cup Z| \ge (1-\alpha)^2 n$. Renumber the X_i if necessary so that either $\mathrm{sign}(X_1) = \mathrm{sign}(X_2)$ or $\mathrm{sign}(X_1) = 0$, and so that $|\mathrm{wt}(X_1)| \le |\mathrm{wt}(X_2)|$. Let $A \leftarrow X_1$, $B \leftarrow X_2$, $C \leftarrow Z$, and $D \leftarrow X_3$.

(2) Additional invariants: $\mathrm{wt}(A) \cdot \mathrm{wt}(B) \ge 0$, and $|\mathrm{wt}(A)| \le |\mathrm{wt}(B)|$.

(3) Use the original n^λ-separator algorithm to separate at each iteration.

(4) To choose a fragment to put into a bucket, first exchange X and Y if necessary to make $|\mathrm{wt}(X)| \le |\mathrm{wt}(Y)|$. Then take one of the following three actions:

Case 1 (Y fits into B without sign change).
If $\mathrm{sign}(Y) = \mathrm{sign}(B)$, then let $B \leftarrow B \cup Y$;

Case 2 (X fits into B without sign change).
Otherwise, if $\mathrm{sign}(X) = \mathrm{sign}(B)$ or if $|\mathrm{wt}(X)| \le |\mathrm{wt}(B)|$ or if $\mathrm{sign}(A) = \mathrm{sign}(B) = 0$, then let $B \leftarrow B \cup X$;

Case 3 (Both A and B change sign).
In the remaining case, $|\mathrm{wt}(A)| \le |\mathrm{wt}(B)| < |\mathrm{wt}(X)| \le |\mathrm{wt}(Y)|$, and $\mathrm{sign}(Y) = \mathrm{sign}(X) = -\mathrm{sign}(B)$. It is still true that either $\mathrm{sign}(A) = \mathrm{sign}(B)$ or $\mathrm{sign}(A) = 0$. Let $A \leftarrow A \cup X$ and $B \leftarrow B \cup Y$;

After putting a fragment in a bucket, exchange A and B if necessary to make $|\mathrm{wt}(A)| \leq |\mathrm{wt}(B)|$;

It is straightforward to verify that the three cases in Step 4 are exhaustive, and that the modified invariant is preserved. At the end, $\mathrm{wt}(A) + \mathrm{wt}(B) + \mathrm{wt}(C) = 0$, and $\mathrm{wt}(A)$ and $\mathrm{wt}(B)$ do not have opposite signs, so $|\mathrm{wt}(A)| + |\mathrm{wt}(B)| = |\mathrm{wt}(C)|$. This implies $|\mathrm{wt}(A)| \leq |\mathrm{wt}(C)|$ and $|\mathrm{wt}(B)| \leq |\mathrm{wt}(C)|$ as desired.

The three-way separation in Step 1 ensures that both A and B begin with enough vertices that $|A| \leq \gamma n$ and $|B| \leq \gamma n$ at the end, where $\gamma = 1 - (1-\alpha)^2 < 1$. This completes the proof. $\square$

This algorithm has two key ideas. The first is the initial two separations, which ensure that the fraction of total size in each bucket will be bounded away from 1. The second is the invariant that the total weights of the two buckets always have the same sign (or one is zero). We can maintain the invariant because either some fragment can be added to some bucket without changing either bucket's sign, or else one fragment can be added to each bucket, changing both buckets' signs. Both of these ideas are used in slightly messier form in the next section.

3 Balancing by modifying weights

The purpose of this section is to prove a technical lemma. While using *generic-separate* to partition a graph with several kinds of weight on its vertices, the lemma describes a way to modify the weights at each iteration. According to the modified weights, the separator has no weight at all, and within each bucket all the kinds of weight are equal. At the end of the separation, then, all the kinds of modified weight are separated in the same ratio, though that ratio may not be close to 1:1. The modifications are done in a way that preserves the sum of each kind of weight, and that can only increase the weights of vertices outside the separator.

All of the weights in this section are nonnegative. The next section will use the separator result from Lemma 4 in another application of *generic-separate* to prove the main results of the paper.

As in Theorem 3, the algorithm begins with two separations to guarantee that both buckets have a fraction bounded away from 1 of the vertices, and it preserves an invariant saying that the buckets do not have opposite sign (in this case, the relevant sign is that of the difference between two kinds of weight). At the end, the ratios for the various weights are balanced exactly by moving weight (but not vertices) out of the separator into the parts. The details of this final modification are tedious, but simple.

Lemma 4 *Let G be an n-vertex graph whose subgraphs all satisfy an n^λ-separator theorem with parameters $\alpha < 1$ and $\beta > 0$. Suppose the vertices of G have k unrelated nonnegative weights $\mathrm{wt}_i(v)$ (for $1 \le i \le k$), each summing to one: $\sum_v \mathrm{wt}_i(v) = 1$. Then there exist a partition (A, B, C) of the vertices of G and new nonnegative weight functions $\mathrm{wt}'_i(v)$ such that, for $1 \le i \le k$,*

- $\mathrm{wt}'_i(v) \ge \mathrm{wt}_i(v)$ *for all vertices $v \in A \cup B$.*
- $\mathrm{wt}'_i(A \cup B \cup C) = 1$.
- $\mathrm{wt}'_1(A) = \cdots = \mathrm{wt}'_k(A)$ *and* $\mathrm{wt}'_1(B) = \cdots = \mathrm{wt}'_k(B)$.
- $\mathrm{wt}'_1(C) = \cdots = \mathrm{wt}'_k(C) = 0$ *unless $A \cup B$ is empty, in which case* $\mathrm{wt}'_i(v) = \mathrm{wt}_i(v)$ *for all i and v.*
- *No edge joins A and B.*
- $|A| \le \alpha_k n$, $|B| \le \alpha_k n$, *and* $|C| \le \beta_k n^\lambda$.

Here α_k and β_k are constants that depend only on α, β, λ, and k. If the original n^λ-separator algorithm runs in $t(n)$ time, the partition and the modified weights can be found in $O(t(n))$ time.

Proof: We prove the lemma by induction on the number k of weights. If $k = 0$, the lemma is just the original n^λ-separator theorem, so we can take $\alpha_0 = \alpha$ and $\beta_0 = \beta$.

The inductive step is much like the proof of Theorem 3, except that $(\mathrm{wt}'_k - \mathrm{wt}'_1)$ plays the role of the single mixed-sign weight in that proof, and the separator subroutine called in Steps 1 and 3 is this one, called recursively with the first $k-1$ weights. The details follow. We take $k \ge 1$ and assume inductively that the lemma holds for $k-1$ weights.

We construct an algorithm *multisep*$_k$ by filling in the details of algorithm *generic-separate*. Notice that the weights wt'_i change during the recursive calls to *multisep*$_{k-1}$, and they are also changed explicitly at the end of the algorithm. The conclusions about *generic-separate* in Lemma 1 do not depend on the weights, so the fact that the weights change as we go along does not affect those conclusions.

(1) Initially, let $\mathrm{wt}'_i(v) \leftarrow \mathrm{wt}_i(v)$ for all i and v. Using the modified weights $\mathrm{wt}'_1, \ldots, \mathrm{wt}'_{k-1}$, call subroutine $\mathit{multisep}_{k-1}$ once on G and then once again on the larger of the resulting pieces.

The result is a partition of the vertices of G into a separator Z with at most $2\beta_{k-1}n^\lambda$ vertices, and three pieces X_1, X_2, and X_3 such that no edge joins X_i to X_j if $i \neq j$, and such that $|X_i \cup Z| \geq (1-\alpha_{k-1})^2 n$. In addition, $\mathrm{wt}'_1(X_i) = \cdots = \mathrm{wt}'_{k-1}(X_i)$ for $i = 1, 2, 3$.

We can renumber the X_i so that

$$(\mathrm{wt}'_k(X_1) - \mathrm{wt}'_1(X_1)) \cdot (\mathrm{wt}'_k(X_2) - \mathrm{wt}'_1(X_2)) \geq 0,$$

$$|\,\mathrm{wt}'_k(X_1) - \mathrm{wt}'_1(X_1)| \leq |\,\mathrm{wt}'_k(X_2) - \mathrm{wt}'_1(X_2)|.$$

Do so, and let $A \leftarrow X_1$, $B \leftarrow X_2$, $C \leftarrow Z$, and $D \leftarrow X_3$.

(2) Additional invariants:

$$(\mathrm{wt}'_k(A) - \mathrm{wt}'_1(A)) \cdot (\mathrm{wt}'_k(B) - \mathrm{wt}'_1(B)) \geq 0,$$

$$|\,\mathrm{wt}'_k(A) - \mathrm{wt}'_1(A)| \leq |\,\mathrm{wt}'_k(B) - \mathrm{wt}'_1(B)|,$$

$$\mathrm{wt}'_1(A) = \cdots = \mathrm{wt}'_{k-1}(A),$$

$$\mathrm{wt}'_1(B) = \cdots = \mathrm{wt}'_{k-1}(B),$$

$$\mathrm{wt}'_1(C) = \cdots = \mathrm{wt}'_{k-1}(C) = 0,$$

$$\mathrm{wt}'_i(v) \geq \mathrm{wt}_i(v) \quad \text{for all } i \text{ and all } v \in A \cup B \cup D,$$

$$\sum_v \mathrm{wt}'_i(v) = 1 \quad \text{for all } i.$$

(3) At each iteration, separate the remaining vertices D by a recursive call to $\mathit{multisep}_{k-1}$ using weights $\mathrm{wt}'_1, \ldots, \mathrm{wt}'_{k-1}$. This changes the weights wt'_i but preserves the invariants.

(4) Choose one of X or Y to place into one of A or B. Any choice preserves all the invariants except the first two "additional invariants." To preserve them, proceed as follows. Exchange X and Y if necessary to make $|\,\mathrm{wt}'_k(X) - \mathrm{wt}'_1(X)| \leq |\,\mathrm{wt}'_k(Y) - \mathrm{wt}'_1(Y)|$. Take one of the following three actions:

Case 1 (Y fits into B without sign change).
If $\mathrm{sign}(\mathrm{wt}'_k(Y) - \mathrm{wt}'_1(Y)) = \mathrm{sign}(\mathrm{wt}'_k(B) - \mathrm{wt}'_1(B))$, then let $B \leftarrow B \cup Y$;

Case 2 (X fits into B without sign change).
Otherwise, if any of the following is true:

$$\operatorname{sign}(\mathrm{wt}'_k(X) - \mathrm{wt}'_1(X)) = \operatorname{sign}(\mathrm{wt}'_k(B) - \mathrm{wt}'_1(B)),$$

$$|\,\mathrm{wt}'_k(X) - \mathrm{wt}'_1(X)| \leq |\,\mathrm{wt}'_k(B) - \mathrm{wt}'_1(B)|,$$

$$\mathrm{wt}'_k(A) - \mathrm{wt}'_1(A) = \mathrm{wt}'_k(B) - \mathrm{wt}'_1(B) = 0,$$

then let $B \leftarrow B \cup X$;

Case 3 (Both A and B change sign).
In the remaining case, the difference $\mathrm{wt}'_k - \mathrm{wt}'_1$ has the opposite sign in both X and Y from that in B, and has larger magnitude in both X and Y than in either A or B. Therefore if we put one of X and Y into each of A and B, the sign of the difference will change in both A and B, preserving the first invariant. Thus let $A \leftarrow A \cup X$ and $B \leftarrow B \cup Y$;

To preserve the second invariant, exchange A and B if necessary to make $|\,\mathrm{wt}'_k(A) - \mathrm{wt}'_1(A)| \leq |\,\mathrm{wt}'_k(B) - \mathrm{wt}'_1(B)|$.

After the final iteration of the **while** loop, set D is empty and sets A, B, and C partition the vertices. We have $\mathrm{wt}'_1(A) = \cdots = \mathrm{wt}'_{k-1}(A)$, $\mathrm{wt}'_1(B) = \cdots = \mathrm{wt}'_{k-1}(B)$, and $\mathrm{wt}'_1(C) = \cdots = \mathrm{wt}'_{k-1}(C)$. It is possible that C may have nonzero weight of the first $k-1$ kinds, because the recursive call in the final iteration may have returned $X = Y = \emptyset$ and perforce left all the remaining weight in Z.

At this point we perform a final modification to make $\mathrm{wt}'_i(C) = 0$ for all i, and to ensure that wt'_k is divided in the same ratio as the other weights. The idea is to move all the remaining weight, of all kinds, out of C into A and B, appealing to the fact that (by the first invariant) the sign of $\mathrm{wt}'_k - \mathrm{wt}'_1$ is the same in both A and B to argue that there is enough weight of the right kinds in C to make all the modified weights equal in each of A and B. Here are the details.

If $A = B = \emptyset$ we are done. Otherwise we move all the weights out of C as follows.

First, consider the case that $\mathrm{wt}'_1(C) = \cdots = \mathrm{wt}'_{k-1}(C) = 0$ (see Figure 3). Then

$$\mathrm{wt}'_k(A) \leq \mathrm{wt}'_1(A) \quad \text{and} \quad \mathrm{wt}'_k(B) \leq \mathrm{wt}'_1(B)$$

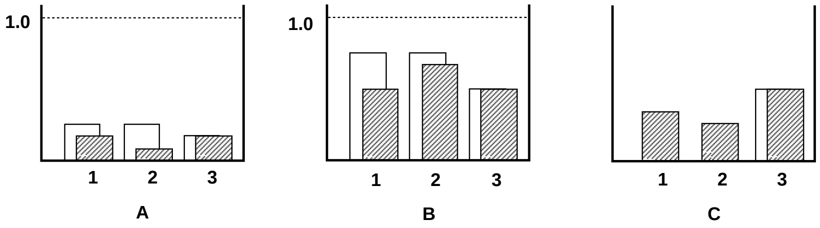

Figure 3: Before final modification (first case). Shading is wt_i; white is wt'_i.

and (since each kind of weight sums to 1)

$$\mathrm{wt}'_k(C) = (\mathrm{wt}'_1(A) - \mathrm{wt}'_k(A)) + (\mathrm{wt}'_1(B) - \mathrm{wt}'_k(B)).$$

Choose arbitrary vertices $a_0 \in A$ and $b_0 \in B$ (if any). Increase $\mathrm{wt}'_k(a_0)$ by $\mathrm{wt}'_1(A) - \mathrm{wt}'_k(A)$, and increase $\mathrm{wt}'_k(b_0)$ by $\mathrm{wt}'_1(B) - \mathrm{wt}'_k(B)$. Set $\mathrm{wt}'_k(v) \leftarrow 0$ for every $v \in C$. The result is $\mathrm{wt}'_1(A) = \cdots = \mathrm{wt}'_k(A)$, $\mathrm{wt}'_1(B) = \cdots = \mathrm{wt}'_k(B)$, and $\mathrm{wt}'_1(C) = \cdots = \mathrm{wt}'_k(C) = 0$. (If B is empty we cannot choose such a b_0, but then $\mathrm{wt}'_1(B) - \mathrm{wt}'_k(B) = 0$ so we would be moving 0 weight into B in any case.)

Second, consider the case that $\mathrm{wt}'_1(C) = \cdots = \mathrm{wt}'_{k-1}(C) \neq 0$, and $\mathrm{wt}'_k(C) \geq \mathrm{wt}'_1(C)$. The first invariant and the fact that each weight sums to one imply that

$$\mathrm{wt}'_k(A) \leq \mathrm{wt}'_1(A) \quad \text{and} \quad \mathrm{wt}'_k(B) \leq \mathrm{wt}'_1(B).$$

Choose an arbitrary vertex $v_0 \in A \cup B$. For each $1 \leq i < k$, add $\mathrm{wt}'_i(C)$ to $\mathrm{wt}'_i(v_0)$, and set $\mathrm{wt}'_i(v) \leftarrow 0$ for every $v \in C$. This preserves the sums of weights and the facts that $\mathrm{wt}'_1(A) = \cdots = \mathrm{wt}'_{k-1}(A)$ and $\mathrm{wt}'_1(B) = \cdots = \mathrm{wt}'_{k-1}(B)$, and it makes $\mathrm{wt}'_1(C) = \cdots = \mathrm{wt}'_{k-1}(C) = 0$. It also preserves the facts that $\mathrm{wt}'_k(A) \leq \mathrm{wt}'_1(A)$ and $\mathrm{wt}'_k(B) \leq \mathrm{wt}'_1(B)$. Now proceed as in the first case above to move wt'_k out of C.

The third and final case is that $\mathrm{wt}'_1(C) = \cdots = \mathrm{wt}'_{k-1}(C) \neq 0$, and $\mathrm{wt}'_k(C) < \mathrm{wt}'_1(C)$. This is just like the second case, but the roles of wt'_k and $\mathrm{wt}'_1 = \cdots = \mathrm{wt}'_{k-1}$ are reversed: choose $v_0 \in A \cup B$, add $\mathrm{wt}'_k(C)$ to $\mathrm{wt}'_k(v_0)$, and set $\mathrm{wt}'_k(v) \leftarrow 0$ for every $v \in C$. Again, this preserves the sums of weights and the facts that $\mathrm{wt}'_1(A) = \cdots = \mathrm{wt}'_{k-1}(A)$ and $\mathrm{wt}'_1(B) = \cdots = \mathrm{wt}'_{k-1}(B)$, and this time it makes $\mathrm{wt}'_k(C) = 0$. It also preserves the facts that $\mathrm{wt}'_1(A) \leq \mathrm{wt}'_k(A)$ and $\mathrm{wt}'_1(B) \leq \mathrm{wt}'_k(B)$. Now proceed as in the first case above, reversing the roles of wt'_k and $\mathrm{wt}'_1 = \cdots = \mathrm{wt}'_{k-1}$, to move $\mathrm{wt}'_1 = \cdots = \mathrm{wt}'_{k-1}$ out of C.

This finishes the algorithm. All the modifications preserved sums of weights, so the modified weights all sum to 1. The weights of vertices outside the separator were only increased, so $\mathrm{wt}'_i(v) \geq \mathrm{wt}_i(v)$ for $v \in A \cup B$.

Just as in Theorem 3, the three-way separation in Step (1) ensures that both A and B begin with enough vertices so that $|A| \leq \alpha_k n$ and $|B| \leq \alpha_k n$ at the end, for $\alpha_k = 1 - (1 - \alpha_{k-1})^2 < 1$. The new constant in the separator size is $\beta_k \leq \beta_{k-1}(1 + 1/(1 - \alpha_{k-1}))$. $\square$

4 Multiple-weight separator theorems

In this section we use Lemma 4 to prove an exact separator theorem for multiple nonnegative weights, and then use that theorem to prove a separator theorem for multiple arbitrary weights.

Theorem 5 (Exact multiple-weight separator theorem)
Let G be an n-vertex graph whose subgraphs all satisfy an n^λ-separator theorem with parameters $\alpha < 1$ and $\beta > 0$. Suppose the vertices of G have k nonnegative weights $\mathrm{wt}_i(v)$, summing to s_i, for $1 \leq i \leq k$. Then for any $t_A > 0$ and $t_B > 0$ such that $t_A + t_B = 1$, the vertices of G can be partitioned into sets A, B, and C such that

- $\mathrm{wt}_i(A) \leq t_A s_i \quad$ for $1 \leq i \leq k$.
- $\mathrm{wt}_i(B) \leq t_B s_i \quad$ for $1 \leq i \leq k$.
- *No edge joins A and B.*
- $|C| \leq \delta n^\lambda$.

Here δ is a constant that depends only on α, β, λ, and k. If the original n^λ-separator algorithm runs in $t(n)$ time, the partition can be found in $O(t(n))$ time.

Proof: Without loss of generality we can ignore any kind of weight that is identically 0, so that $s_i > 0$ for all i, and then assume that the weights have been scaled so that $s_i = 1$ for all i.

The proof uses algorithm *generic-separate*. In place of the call to *separate*, we call the algorithm of Lemma 4. That algorithm returns new weights for some of the vertices, which we use in the next call. We will use $\mathrm{wt}_i(v)$ to refer to the original weights, which never change, and we will use $\mathrm{wt}'_i(v)$ to refer to the weights that change at each iteration. The constants in Lemma 4 do not depend on the weights, so the fact that the weights change as we go along does not affect the conclusions of Lemma 1. We fill in the details of *generic-separate* as follows.

(1) Initially, let $\mathrm{wt}'_i(v) \leftarrow \mathrm{wt}_i(v)$ for all v and i;

(2) Additional invariants:

$$\mathrm{wt}'_i(v) \geq \mathrm{wt}_i(v) \text{ for all } v \in A \cup B \cup D \text{ and for all } i;$$
$$t_A - \mathrm{wt}'_1(A) \geq t_B - \mathrm{wt}'_1(B) \geq 0;$$
$$\mathrm{wt}'_i(G) = 1 \text{ for all } i;$$
$$\mathrm{wt}'_i(C) = 0 \text{ for all } i;$$
$$\mathrm{wt}'_1(A) = \cdots = \mathrm{wt}'_k(A);$$
$$\mathrm{wt}'_1(B) = \cdots = \mathrm{wt}'_k(B);$$

(3) To separate at each iteration, use the algorithm from Lemma 4, starting with weights wt'_i. The algorithm returns new weights wt'_i for the vertices of D, and a partition X, Y, Z of D satisfying the lemma.

(4) To choose a fragment to put into a bucket:
exchange X and Y if necessary to make $\mathrm{wt}'_1(X) \leq \mathrm{wt}'_1(Y)$;
$A \leftarrow A \cup X$;
exchange A and B if necessary to make
$t_A - \mathrm{wt}'_1(A) \geq t_B - \mathrm{wt}'_1(B)$;

Figure 4 shows an iteration of the algorithm. The invariant implies that the weights wt'_i in the call to the separator algorithm in Step 3 satisfy the hypotheses of Lemma 4. It is straightforward to check that Step 4 preserves the invariant, using the conclusions of Lemma 4 and noting that (before the modification $A \leftarrow A \cup X$ is made)

$$t_A - (\mathrm{wt}'_1(A) + \mathrm{wt}'_1(X)) \geq t_B - (\mathrm{wt}'_1(B) + \mathrm{wt}'_1(Y)),$$
$$\mathrm{wt}'_1(A) + \mathrm{wt}'_1(X) + \mathrm{wt}'_1(B) + \mathrm{wt}'_1(Y) \leq 1.$$

At the end, the invariant implies that the original weights of A and B satisfy the inequalities of the theorem and, since D is empty, A, B, and C partition the vertices of G. $\Box$

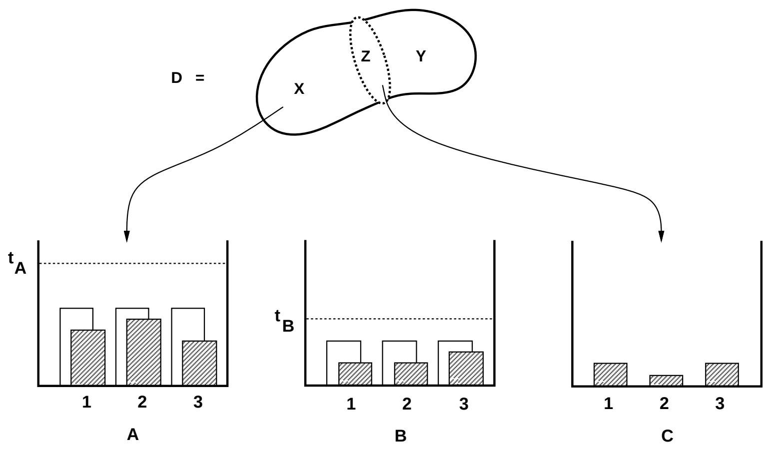

Figure 4: One iteration of Theorem 5. Shading is wt_i; white is wt'_i.

Gilbert and Tarjan [9] used a different technique to prove that several kinds of nonnegative weight could be simultaneously divided, with each part having at most $\frac{1}{2}$ of the weight of one kind and $\frac{1}{2} + \epsilon$ of the weight of the other kinds. The separator size was $O(n^\lambda/\epsilon)$. Theorem 5 is stronger because the precision of the split and the separator size do not depend on each other; also, the weights can be divided in any desired ratio.

We now state several corollaries of Theorem 5. The first just says that any number of nonnegative weights can be divided exactly in half by the same separator.

Corollary 6 (Exact multiple-weight 50-50 separator theorem) *Let G be an n-vertex graph whose subgraphs all satisfy an n^λ-separator theorem with parameters $\alpha < 1$ and $\beta > 0$. Suppose the vertices of G have k unrelated nonnegative sets of weights $\mathrm{wt}_i(v)$ summing to s_i, for $1 \le i \le k$. Then the vertices of G can be partitioned into sets A, B, and C such that*

- $\mathrm{wt}_i(A) \le s_i/2 \quad \text{for } 1 \le i \le k.$
- *No edge joins A and B.*
- $|C| \le \delta n^\lambda.$

Here δ is a constant that depends only on α, β, λ, and k. If n^λ-separators can be found in $t(n)$ time, this partition can be found in $O(t(n))$ time.

Proof: Take $t_A = t_B = \frac{1}{2}$ in Theorem 5. $\Box$

Our most general result says that we can divide any number of weights of arbitrary sign evenly by the same separator. Again this raises the question of what it means to divide a weight "evenly" when its values may be both positive and negative, and its sum over the whole graph may be any real number. One definition of an "even division" into two parts A and B by a separator C is

$$|\operatorname{wt}(A) - \operatorname{wt}(B)| \leq \sum_{v \in C} |\operatorname{wt}(v)|.$$

If the weights are all positive and sum to s, this implies that $\operatorname{wt}(A)$ and $\operatorname{wt}(B)$ are at most $s/2$. If the weights sum to 0, this condition is slightly weaker than Theorem 5's condition that $|\operatorname{wt}(A)|$ and $|\operatorname{wt}(B)|$ are at most $|\operatorname{wt}(C)|$. The next corollary says that a separator theorem for vertices implies a separator theorem, by this criterion, for any number of weights.

Corollary 7 (Multiple mixed-weight separator theorem)
Let G be an n-vertex graph whose subgraphs all satisfy an n^λ-separator theorem with parameters $\alpha < 1$ and $\beta > 0$. Suppose the vertices of G have k unrelated sets of weights $\operatorname{wt}_i(v)$ for $1 \leq i \leq k$, which can be arbitrary real numbers. Then the vertices of G can be partitioned into sets A, B, and C such that

- $|\operatorname{wt}_i(A) - \operatorname{wt}_i(B)| \leq \sum_{v \in C} |\operatorname{wt}_i(v)| \quad$ for $1 \leq i \leq k$.
- *No edge joins A and B.*
- $|C| \leq \delta n^\lambda$.

Here δ is a constant that depends only on α, β, λ, and k. If n^λ-separators can be found in $t(n)$ time, this partition can be found in $O(t(n))$ time.

Proof: Let $\operatorname{wt}_i^+(v) = \max(\operatorname{wt}_i(v), 0)$ and let $\operatorname{wt}_i^-(v) = \max(-\operatorname{wt}_i(v), 0)$, for all i and v. Use Corollary 6 to partition the vertices into a separator C and two pieces A and B with $\operatorname{wt}_i^+(A) \leq \sum_v \operatorname{wt}_i^+(v)/2$, $\operatorname{wt}_i^-(A) \leq \sum_v \operatorname{wt}_i^-(v)/2$, $\operatorname{wt}_i^+(B) \leq \sum_v \operatorname{wt}_i^+(v)/2$, and $\operatorname{wt}_i^-(B) \leq \sum_v \operatorname{wt}_i^-(v)/2$, for all i. Then $|\operatorname{wt}_i(A) - \operatorname{wt}_i(B)| \leq |\operatorname{wt}_i^+(A) - \operatorname{wt}_i^+(B)| + |\operatorname{wt}_i^-(A) - \operatorname{wt}_i^-(B)| \leq \operatorname{wt}_i^+(C) + \operatorname{wt}_i^-(C) = \sum_{v \in C} |\operatorname{wt}(v)|$. $\Box$

5 Remarks

For graphs embedded in surfaces, we can consider placing weights on the edges and faces of the embedding as well as on the vertices, and requiring the separator to be a set of edges and vertices whose removal from the embedded graph disconnects the surface. For example, Miller [18] showed that with nonnegative weight on the vertices and faces of an embedding, a triangulated planar graph has a separator that is a simple cycle of $O(n^{1/2})$ vertices. We expect that our results extend to the case of negative and multiple edge and face weights as well.

However, the separators in our theorems cannot necessarily be chosen to be single simple cycles. Consider an n-vertex planar graph obtained by triangulating the surface and ends of a solid cylinder of height $\Theta(n)$ and circumference 1 with triangles of about the same size. Let $n^{2/3}$ vertices near one end each have weight 1, let $n^{2/3}$ vertices near the opposite end each have weight -1, and let the remaining vertices have weight 0. The separator in Theorem 3 cannot include all the positive or all the negative vertices, so it must separate both the positive and negative vertices; but a single cycle that could do so would have to be $\Omega(n)$ vertices long. Similar examples show that Theorem 5 and its corollaries do not hold for single cycle separators in triangulated graphs.

We conclude with two open questions. First, is it possible to strengthen Corollary 7 by replacing the condition $|\,\mathrm{wt}_i(A) - \mathrm{wt}_i(B)| \leq \sum_{v \in C} |\,\mathrm{wt}_i(v)|$ by $|\,\mathrm{wt}_i(A) - \mathrm{wt}_i(B)| \leq |\,\mathrm{wt}_i(C)|$ as in Theorem 3? We conjecture that the answer is yes. (However, we note that this does not change the result for nonnegative weights; and another open question is whether there is any application for negative vertex weights!) Second, the constants in all the results worsen rapidly as the number of kinds of weight increases; in Lemma 4, β_k and $1/(1-\alpha_k)$ grow double-exponentially with k. Is it possible to improve them?

References

[1] Noga Alon, Paul Seymour, and Robin Thomas. A separator theorem for graphs with an excluded minor and applications. In *Proceedings of the 22th Annual Symposium on Theory of Computing*. ACM, 1990.

[2] Hristo N. Djidjev. A separator theorem. *Comptes Rendus de l' Academie Bulgare des Sciences*, 34:643–645, 1981.

[3] Hristo N. Djidjev and John R. Gilbert. Separators in graphs with negative or multiple vertex weights. Technical Report CSL 90–8, Xerox Palo Alto Research Center, 1990.

[4] Zvi Galil, Ravi Kannan, and Endre Szemeredi. On nontrivial separators for k-page graphs and simulations by nondeterministic one-tape Turing machines. In *Proceedings of the 18th Annual Symposium on Theory of Computing*, pages 39–49. ACM, 1986.

[5] John R. Gilbert, Joan P. Hutchinson, and Robert Endre Tarjan. A separator theorem for graphs of bounded genus. *Journal of Algorithms*, 5:391–407, 1984.

[6] John R. Gilbert, Donald J. Rose, and Anders Edenbrandt. A separator theorem for chordal graphs. *SIAM Journal on Algebraic and Discrete Methods*, 5:306–313, 1984.

[7] John R. Gilbert and Robert Endre Tarjan. The analysis of a nested dissection algorithm. *Numerische Mathematik*, 50:377–404, 1987.

[8] John R. Gilbert and Earl Zmijewski. A parallel graph partitioning algorithm for a message-passing multiprocessor. *International Journal of Parallel Programming*, 16:427–449, 1987.

[9] John Russell Gilbert. *Graph Separator Theorems and Sparse Gaussian Elimination.* PhD thesis, Stanford University, 1980.

[10] D. Hoey and C. E. Leiserson. A layout for the shuffle-exchange graph. Technical Report CMU–CS–80–139, Carnegie-Mellon University, 1980.

[11] Camille Jordan. Sur les assemblages de lignes. *Journal Reine Angew. Math.*, 70:185–190, 1869.

[12] F. Thomas Leighton. New lower bound techniques for vlsi. In *Proceedings of the 22nd Annual Symposium on Foundations of Computer Science*, pages 1–12. IEEE, 1981.

[13] Charles E. Leiserson. Area-efficient graph layouts (for VLSI). In *Proceedings of the 21st Annual Symposium on Foundations of Computer Science*, pages 270–281, 1980.

[14] Ming Li. Simulating two pushdowns by one tape in $O(n^{1.5}\sqrt{\log n})$ time. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science*, pages 383–393. IEEE, 1985.

[15] Richard J. Lipton, Donald J. Rose, and Robert Endre Tarjan. Generalized nested dissection. *SIAM Journal on Numerical Analysis*, 16:346–358, 1979.

[16] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36:177–189, 1979.

[17] Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. *SIAM Journal on Computing*, 9:615–627, 1980.

[18] Gary L. Miller. Finding small simple cycle separators for 2-connected planar graphs. In *Proceedings of the 16th Annual Symposium on Theory of Computing*, pages 376–382. ACM, 1984.

[19] Gary L. Miller, Shanghua Teng, and Steven A. Vavasis. A unified geometric approach to graph separators. In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science*, pages 538–547. IEEE, 1991.

[20] Gary L. Miller and William Thurston. Separators in two and three dimensions. In *Proceedings of the 22nd Annual Symposium on Theory of Computing*. ACM, 1990.

[21] Gary L. Miller and Steven A. Vavasis. Density graphs and separators. In *Proceedings of the Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 331–336. SIAM, 1991.

[22] Shang-Hua Teng. *Points, Spheres, and Separators: A Unified Geometric Approach to Graph Partitioning*. PhD thesis, Carnegie-Mellon University, Pittsburgh, Pennsylvania, August 1991.

[23] Stephen A. Vavasis. Automatic domain partitioning in three dimensions. To appear in the *SIAM Journal on Scientific and Statistical Computing*.