

Using PageRank to Locally Partition a Graph

Reid Andersen, Fan Chung, and Kevin Lang

Abstract. A local graph partitioning algorithm finds a cut near a specified starting vertex, with a running time that depends largely on the size of the small side of the cut, rather than the size of the input graph. In this paper, we present a local partitioning algorithm using a variation of PageRank with a specified starting distribution. We derive a mixing result for PageRank vectors similar to that for random walks, and we show that the ordering of the vertices produced by a PageRank vector reveals a cut with small conductance. In particular, we show that for any set C with conductance Φ and volume k , a PageRank vector with a certain starting distribution can be used to produce a set with conductance $O(\sqrt{\Phi \log k})$. We present an improved algorithm for computing approximate PageRank vectors, which allows us to find such a set in time proportional to its size. In particular, we can find a cut with conductance at most ϕ , whose small side has volume at least 2^b , in time $O(2^b \log^2 m / \phi^2)$ where m is the number of edges in the graph. By combining small sets found by this local partitioning algorithm, we obtain a cut with conductance ϕ and approximately optimal balance in time $O(m \log^4 m / \phi^2)$.

I. Introduction

One of the central problems in algorithmic design is the problem of finding a cut where the ratio between the number of edges crossing the cut and the size of the smaller side of the cut is small. There is a large literature of research papers on this topic, with applications in numerous areas.

Partitioning algorithms that find such cuts can be applied recursively to solve more complicated problems, including finding balanced cuts, k -way partitions,

and hierarchical clusterings [Borgs et al. 04, Kannan et al. 04, Leighton and Rao 88, Simon and Teng 97, Spielman and Teng 96]. The running time of these recursive algorithms can be large if the cuts found at each step are unbalanced. This is particularly evident when applying spectral partitioning, which produces a cut with approximately optimal conductance, but with no guarantee on the balance.

Recently, Spielman and Teng addressed this problem by introducing a local partitioning algorithm called Nibble, which finds a small cut near a specified starting vertex in time proportional to the size of the small side of the cut. The small cuts found by Nibble can be combined to form balanced cuts in nearly linear time, and the resulting balanced cut algorithm, Partition, is used as a subroutine for finding multiway partitions, sparsifying graphs, and solving diagonally dominant linear systems [Spielman and Teng 04]. The analysis of the Nibble algorithm is based on a mixing result by Lovász and Simonovits [Lovász and Simonovits 90, Lovász and Simonovits 93], which shows that a cut with small conductance can be found by simulating a random walk starting from a single vertex for sufficiently many steps.

In this paper, we present a local graph partitioning algorithm that finds cuts by computing and examining PageRank vectors. A PageRank vector is a weighted sum of the probability distributions obtained by taking a sequence of random walk steps starting from a specified initial distribution. The weight placed on the distribution obtained after t walk steps decreases exponentially in t , with the rate of decay determined by a parameter called the *teleport probability*. A PageRank vector can also be viewed as the solution of a system of linear equations, which we will describe in more detail in Section 2. Each of the PageRank vectors we compute has its starting distribution on a single starting vertex, and we prove that a PageRank vector produces a cut which approximates the best cut near its starting vertex. This cut can be found by performing a *sweep* over the PageRank vector, which involves examining the vertices of the graph in an order determined by the PageRank vector, and computing the conductance of each set produced by this order. The analysis of our algorithm is based on the following results:

- We give an algorithm for approximating a PageRank vector by another PageRank vector with a slightly different starting distribution, based on a technique introduced by Jeh and Widom [Jeh and Widom 03]. This allows us to compute a PageRank vector with teleport probability α and with an amount of error sufficiently small for finding a cut with volume k , in time $O(k/\alpha)$. (The volume of a set S is defined to be the sum of the degrees over all vertices in S , and the volume of a cut is the minimum of the volumes of the two sides.)

- We prove a mixing result for PageRank vectors, which shows that if a PageRank vector with teleport probability α has significantly more probability than the stationary distribution on some set of vertices with volume k , a sweep over that PageRank vector will produce a cut with conductance $O(\sqrt{\alpha \log k})$.
- We show that for any set C , and for many vertices v contained in C , a PageRank vector whose starting vertex is v , and whose teleport probability is greater than the conductance of C , has a significant fraction of its probability contained in C .

Combining these results yields a partitioning result for PageRank vectors: for any set C with conductance Φ , there are a significant number of starting vertices within C for which a sweep over an appropriate PageRank vector finds a cut with conductance $O(\sqrt{\Phi \log k})$, where k is the volume of C . Such a cut can be found in time $O(k/\Phi + k \log k)$.

To produce balanced cuts in nearly linear time, we must be able to remove a small piece of the graph in time proportional to the volume of that small piece, rather than the volume of C or the volume of the entire graph. We present the PageRank-Nibble algorithm that does this. For many starting vertices within a set C of conductance $O(\phi^2/\log^2 m)$, this algorithm finds a cut with conductance ϕ and volume $O(k)$ in time $O(k \log^2 m/\phi^2)$, provided that we can guess the volume of the smaller side of the cut within a factor of 2. This improves the Nibble algorithm, which runs in time $O(k \log^4 m/\phi^5)$ and requires that C have conductance $O(\phi^3/\log^2 m)$.

We can combine cuts found by PageRank-Nibble into a cut whose conductance is at most ϕ and whose volume is at least half that of any set with conductance $O(\phi^2/\log^2 m)$, in time $O(m \log^4 m/\phi^2)$ where m is the number of edges in the graph. This improves the Partition algorithm, which obtains a cut whose volume is at least half that of any set with conductance $O(\phi^3/\log^2 m)$ in time $O(m \log^6 m/\phi^5)$.

2. Preliminaries

Throughout the paper we will consider a graph G that is undirected and unweighted. This graph has a vertex set $V = \{v_1, \dots, v_n\}$ and an edge set E with m undirected edges. We write $d(v)$ for the degree of vertex v , and let D be the diagonal matrix where $D_{i,i} = d(v_i)$. We let A be the adjacency matrix, where $A_{i,j} = 1$ if and only if there is an edge joining v_i and v_j .

When we consider vectors over the vertex set V , we will write them as row vectors, so the product of a vector p and a matrix A will be written pA . Two vectors we will use frequently are the stationary distribution,

$$\psi_S(x) = \begin{cases} \frac{d(x)}{\text{vol}(S)} & \text{if } x \in S, \\ 0 & \text{otherwise,} \end{cases}$$

and the indicator function,

$$\chi_v(x) = \begin{cases} 1 & \text{if } x = v, \\ 0 & \text{otherwise.} \end{cases}$$

The support $\text{Supp}(p)$ of a vector p is the set of vertices where p is nonzero, and the sum of this vector over a set S of vertices is written as

$$p(S) = \sum_{u \in S} p(u).$$

If the entries of p are positive and $p(V)$ is at most 1, we will refer to $p(S)$ as the amount of probability from p on S .

2.1. PageRank Vectors

The PageRank vector $\text{pr}_\alpha(s)$ is defined to be the unique solution of the linear system

$$\text{pr}_\alpha(s) = \alpha s + (1 - \alpha)\text{pr}_\alpha(s)W. \quad (2.1)$$

Here, α is a constant in $(0, 1]$ called the *teleport probability*, s is a vector called the *starting vector*, and W is the lazy random walk transition matrix $W = \frac{1}{2}(I + D^{-1}A)$. This is superficially different from the standard definition of PageRank, which uses the standard random walk matrix $D^{-1}A$ instead of the lazy random walk matrix W , but the two definitions are equivalent up to a change in α . A proof of this is given in the appendix (Section 8).

PageRank was introduced by Brin and Page [Brin and Page 98, Page et al. 98], who proposed using PageRank with the starting vector $s = \vec{1}/n$ for search ranking. PageRank vectors whose starting vectors are not uniform, but instead represent a combination of topics and web pages, are called *personalized PageRank vectors* and have been used to provide personalized search ranking and context-sensitive search [Berkhin 06, Fogaras and Racz 04, Haveliwala 03, Jeh and Widom 03].

We will consider PageRank vectors where the starting vector is a single vertex. We will also sometimes allow PageRank vectors where the starting vector s has

both positive and negative entries, so we define the positive part of s as follows,

$$s_+(x) = \begin{cases} s(x) & \text{if } s(x) \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Here are some useful properties of PageRank vectors (also see [Haveliwala 03, Jeh and Widom 03]). The proofs are given in the appendix.

Proposition 2.1. *For any starting vector s , and any constant α in $(0, 1]$, there is a unique vector $\text{pr}_\alpha(s)$ satisfying $\text{pr}_\alpha(s) = \alpha s + (1 - \alpha)\text{pr}_\alpha(s)W$.*

Proposition 2.2. *For any fixed value of α in $(0, 1]$, there is a linear transformation R_α such that $\text{pr}_\alpha(s) = sR_\alpha$. Furthermore, R_α is given by the matrix*

$$R_\alpha = \alpha I + \alpha \sum_{t=1}^{\infty} (1 - \alpha)^t W^t. \quad (2.2)$$

This implies that a PageRank vector is a weighted average of lazy random walk vectors,

$$\text{pr}_\alpha(s) = \alpha s + \alpha \sum_{t=1}^{\infty} (1 - \alpha)^t (sW^t). \quad (2.3)$$

This also implies that a PageRank vector $\text{pr}_\alpha(s)$ is linear in its starting vector s .

2.2. Conductance, Sweeps, and Mixing

We recall that the *volume* of a subset $S \subseteq V$ of vertices is

$$\text{vol}(S) = \sum_{x \in S} d(x)$$

and the volume of the entire graph is $\text{vol}(V) = 2m$, where m is the number of edges in the graph. The *edge boundary* of a set is defined to be

$$\partial(S) = \{\{x, y\} \in E \mid x \in S, y \notin S\}.$$

The *conductance* of a set is

$$\Phi(S) = \frac{|\partial(S)|}{\min(\text{vol}(S), 2m - \text{vol}(S))}.$$

A *sweep* is a technique for producing a cut from a vector and is widely used in spectral partitioning [Mihail 89, Spielman and Teng 96]. We will use the following degree-normalized version of a sweep. Given a PageRank vector $p = \text{pr}_\alpha(s)$ with

support size $N_p = |\text{Supp}(p)|$, let v_1, \dots, v_{N_p} be an ordering of the vertices from highest to lowest probability-per-degree, so that

$$\frac{p(v_i)}{d(v_i)} \geq \frac{p(v_{i+1})}{d(v_{i+1})}.$$

This produces a collection of sets, with one set $S_j^p = \{v_1, \dots, v_j\}$ for each integer j in $\{1, \dots, N_p\}$, which we call *sweep sets*. We let $\Phi(p)$ be the smallest conductance of any of these sweep sets:

$$\Phi(p) = \min_{j \in [1, N_p]} \Phi(S_j^p).$$

A cut with conductance $\Phi(p)$ can be found by sorting p/d and computing the conductance of each sweep set. This can be done in time $O(\text{vol}(\text{Supp}(p)) + N_p \log N_p)$.

To measure how a vector p is distributed in the graph, we define a function $p[x]$ that gives an upper bound on the amount of probability on any set of vertices with volume x . We refer to this function as the Lovász-Simonovits curve, since it was introduced by Lovász and Simonovits [Lovász and Simonovits 90, Lovász and Simonovits 93]. This function is defined for all real numbers x in the interval $[0, 2m]$ and is determined by the amount of probability on the sweep sets; we set

$$p[\text{vol}(S_j^p)] = p(S_j^p), \quad \text{for each } j \in [0, n],$$

and define $p[x]$ to be piecewise linear between these points. In other words, for any point $x \in [0, 2m]$, if j is the unique index where x is between $\text{vol}(S_j^p)$ and $\text{vol}(S_{j+1}^p)$, then

$$p[x] = p(S_j^p) + \frac{x - \text{vol}(S_j^p)}{d(v_{j+1})} p(v_{j+1}).$$

The function $p[x]$ is increasing and concave. It is not hard to see that $p[x]$ is an upper bound on the amount of probability from p on any set with volume x ; for any set S , we have

$$p(S) \leq p[\text{vol}(S)].$$

As an example of the notation we will use throughout the paper, the PageRank vector with teleport probability α and starting vector χ_v is written as $\text{pr}_\alpha(\chi_v)$. If we let $p = \text{pr}_\alpha(\chi_v)$, the amount of probability from this PageRank vector on a set S is written as either $p(S)$ or $[\text{pr}_\alpha(\chi_v)](S)$, and the value of the Lovász-Simonovits curve at the point $x = \text{vol}(S)$ is written as $p[\text{vol}(S)]$.

3. Computing Approximate PageRank Vectors

Instead of computing the PageRank vector $\text{pr}_\alpha(s)$ exactly, we will approximate it by another PageRank vector with a slightly different starting vector, $\text{pr}_\alpha(s - r)$, where r is a vector with nonnegative entries. If $r(v) \leq \epsilon d(v)$ for every vertex in the graph, then we say that $\text{pr}_\alpha(s - r)$ is an ϵ -approximate PageRank vector for $\text{pr}_\alpha(s)$.

Definition 3.1. An ϵ -approximate PageRank vector for $\text{pr}_\alpha(s)$ is a PageRank vector $\text{pr}_\alpha(s - r)$ where the vector r is nonnegative and satisfies $r(v) \leq \epsilon d(v)$ for every vertex v in the graph.

We remark that the total difference between an ϵ -approximate PageRank vector $\text{pr}_\alpha(s - r)$ and the PageRank vector $\text{pr}_\alpha(s)$ on a given set S can be bounded in terms of ϵ and $\text{vol}(S)$. Such a bound is given in Section 5.

In this section, we give the $\text{ApproximatePR}(s, \alpha, \epsilon)$ algorithm for computing an ϵ -approximate PageRank vector with small support. The running time of the algorithm depends on ϵ and α , but is independent of the size of the graph.

Theorem 3.2. *The $\text{ApproximatePR}(s, \alpha, \epsilon)$ algorithm has the following properties. For any starting vector s with $\|s\|_1 \leq 1$, and any constant $\epsilon \in (0, 1]$, the algorithm computes an ϵ -approximate PageRank vector p for $\text{pr}_\alpha(s)$. The support of p satisfies $\text{vol}(\text{Supp}(p)) \leq \frac{2}{(1-\alpha)\epsilon}$, and the running time of the algorithm is $O(\frac{1}{\epsilon\alpha})$.*

The proof of Theorem 3.2 is based on a series of facts that we describe below. The starting point is the observation that the PageRank operator commutes with the lazy walk matrix W ,

$$\text{pr}_\alpha(s)W = \text{pr}_\alpha(sW). \quad (3.1)$$

A proof is included in the appendix. By combining Equation (3.1) with the equation that defines PageRank, we derive the following equation:

$$\begin{aligned} \text{pr}_\alpha(s) &= \alpha s + (1 - \alpha)\text{pr}_\alpha(s)W \\ &= \alpha s + (1 - \alpha)\text{pr}_\alpha(sW). \end{aligned} \quad (3.2)$$

Jeh and Widom derived this equation and showed that it provides a flexible way to compute many PageRank vectors simultaneously [Jeh and Widom 03]. A similar approach was used by Berkhin [Berkhin 06]. The algorithms they proposed can be used to compute a single approximate PageRank vector in time $O(\frac{\log n}{\epsilon\alpha})$. The difference of $\log n$ between their running time and ours is the

overhead that they incur by using a heap or priority queue instead of a FIFO queue.

Our algorithm maintains a pair of vectors p and r , starting with the trivial approximation $p = \vec{0}$ and $r = s$, and applies a series of push operations that move probability from r to p while maintaining the invariant $p = \text{pr}_\alpha(s - r)$. Each push operation takes the probability from r at a single vertex u , moves an α fraction of this probability to $p(u)$, and then spreads the remaining $1 - \alpha$ fraction within r by applying a lazy random walk step to the vector $(1 - \alpha)r(u)\chi_u$. We define the push operation more formally in Definition 3.3, and then verify that each push operation does maintain the invariant $p = \text{pr}_\alpha(s - r)$.

Definition 3.3. For the push operation, $\text{push}(u)$, let $p' = p$ and $r' = r$, except for these changes:

1. $p'(u) = p(u) + \alpha r(u)$,
2. $r'(u) = (1 - \alpha)r(u)/2$,
3. for each vertex v such that $(u, v) \in E$, $r'(v) = r(v) + (1 - \alpha)r(u)/(2d(u))$.

Lemma 3.4. Let p' and r' be the result of performing $\text{push}(u)$ on p and r . Then,

$$p = \text{pr}_\alpha(s - r) \quad \implies \quad p' = \text{pr}_\alpha(s - r').$$

Proof. After the push operation, we have

$$\begin{aligned} p' &= p + \alpha r(u)\chi_u, \\ r' &= r - r(u)\chi_u + (1 - \alpha)r(u)\chi_u W. \end{aligned}$$

We will apply Equation (3.2) to $r(u)\chi_u$ to show that $p + \text{pr}_\alpha(r) = p' + \text{pr}_\alpha(r')$. The lemma will follow by the linearity property of PageRank vectors:

$$\begin{aligned} \text{pr}_\alpha(r) &= \text{pr}_\alpha(r - r(u)\chi_u) + \text{pr}_\alpha(r(u)\chi_u) \\ &= \text{pr}_\alpha(r - r(u)\chi_u) + \alpha r(u)\chi_u + \text{pr}_\alpha((1 - \alpha)r(u)\chi_u W) \\ &= \text{pr}_\alpha(r - r(u)\chi_u + (1 - \alpha)r(u)\chi_u W) + \alpha r(u)\chi_u \\ &= \text{pr}_\alpha(r') + p' - p. \end{aligned}$$

□

During each push operation, some probability is moved from r to p , where it remains. Our algorithm performs pushes only on vertices where $r(u) \geq \epsilon d(u)$,

Algorithm 1. (ApproximatePR(s, α, ϵ .)

1. Let $p = \vec{0}$, and $r = s$.
2. While $r(u) \geq \epsilon d(u)$ for some vertex u :
 - (a) Pick any vertex u where $r(u) \geq \epsilon d(u)$.
 - (b) Apply $\text{push}(u)$.
3. Return p and r .

which ensures that a significant amount of probability is moved at each step and allows us to bound the number of pushes required to compute an ϵ -approximate PageRank vector.

Algorithm 1 can be implemented by maintaining a queue containing those vertices u satisfying $r(u) \geq \epsilon d(u)$. At each step, a push operation is performed on the first vertex u in the queue. If $r(u)$ is still at least $\epsilon d(u)$ after the push is performed, then u is placed at the back of the queue: otherwise, u is removed from the queue. If a push operation raises the value of $r(x)$ above $\epsilon d(x)$ for some neighbor x of u , then x is added to the back of the queue. This continues until the queue is empty, at which point all vertices satisfy $r(u) < \epsilon d(u)$. We now show that this algorithm has the properties promised in Theorem 3.2.

Proof of Theorem 3.2. Each push operation preserves the property $p = \text{pr}_\alpha(s - r)$, and the stopping criterion ensures that r satisfies $r(u) < \epsilon d(u)$ at every vertex, so the algorithm returns an ϵ -approximate PageRank vector for $\text{pr}_\alpha(s)$.

To bound the running time, let T be the total number of push operations performed by ApproximatePR, and let d_i be the degree of the vertex where the i th push operation was performed. When the i th push operation was performed, the amount of probability on this vertex was at least ϵd_i , so $\|r\|_1$ decreased by at least $\alpha \epsilon d_i$. Since $\|r\|_1$ was at most 1 initially, we must have $\alpha \epsilon \sum_{i=1}^T d_i \leq 1$, so

$$\sum_{i=1}^T d_i \leq \frac{1}{\epsilon \alpha}. \quad (3.3)$$

It is possible to perform a push operation on the vertex u , and to perform the necessary queue updates, in time proportional to $d(u)$. The running time bound for ApproximatePR follows from Equation (3.3).

To bound the support volume, notice that for each vertex v in $\text{Supp}(p)$, there is some probability remaining on $r(v)$ when the algorithm terminates. In fact,

we must have $r(v) \geq ((1 - \alpha)/2) \cdot \epsilon d(v)$, because when the last push operation was performed at vertex v , $r(v)$ was at least $\epsilon d(v)$, and a $\frac{1-\alpha}{2}$ fraction of that probability remained on $r(v)$. It follows that $\text{vol}(\text{Supp}(p)) \leq \frac{2}{(1-\alpha)\epsilon}$. \square

4. A Mixing Result for PageRank Vectors

In this section, we prove a mixing result for PageRank vectors that is an analogue of the Lovász-Simonovits mixing result for random walks. For any PageRank vector $\text{pr}_\alpha(s)$, we give an upper bound on $\text{pr}_\alpha(s)[x]$ that depends on the smallest conductance $\Phi(\text{pr}_\alpha(s))$ found by performing a sweep over $\text{pr}_\alpha(s)$. We use this mixing result to prove the following theorem, which shows that if there exists a set of vertices S that contains more probability from $\text{pr}_\alpha(s)$ than from the stationary distribution ψ_V by a constant amount, then a sweep over $\text{pr}_\alpha(s)$ finds a cut with conductance $O(\sqrt{\alpha \log(\text{vol}(S))})$.

Theorem 4.1. *If $\text{pr}_\alpha(s)$ is a PageRank vector with $\|s_+\|_1 \leq 1$, and there exists a set S of vertices and a constant δ satisfying*

$$[\text{pr}_\alpha(s)](S) - \psi_V(S) > \delta,$$

then

$$\Phi(\text{pr}_\alpha(s)) < \sqrt{\frac{12\alpha \log(4\sqrt{\text{vol}(S)}/\delta)}{\delta}}.$$

In the remainder of this section, we will derive a sequence of lemmas and a general mixing result that lead to the proof of this theorem. When we eventually apply this theorem, we will apply it to an ϵ -approximate PageRank vector. For the conditions of Theorem 4.1 to hold, ϵ must be small enough that the approximate PageRank vector contains more probability than the stationary distribution on some set of vertices.

The first step toward the proof of Theorem 4.1 is to consider how probability moves when a lazy random step is applied to an arbitrary vector p . To do so, we view each undirected edge $\{u, v\}$ as a pair of directed edges (u, v) and (v, u) . For each directed edge (u, v) we let

$$p(u, v) = \frac{p(u)}{d(u)},$$

and for any set of directed edges A , we define

$$p(A) = \sum_{(u,v) \in A} p(u,v).$$

When a lazy walk step is applied to the vector p , the amount of probability that moves from u to v is $\frac{1}{2}p(u,v)$. For any set S of vertices, we define the set of directed edges into S ,

$$\text{in}(S) = \{(u,v) \in E \mid v \in S\},$$

and the set of directed edges out of S ,

$$\text{out}(S) = \{(u,v) \in E \mid u \in S\}.$$

The following lemma describes the amount of probability on a set S in terms of the amount of probability moving across $\text{in}(S)$ and $\text{out}(S)$.

Lemma 4.2. *For any vector p and any set S of vertices,*

$$pW(S) = \frac{1}{2}(p(\text{in}(S) \cup \text{out}(S)) + p(\text{in}(S) \cap \text{out}(S))).$$

Proof. The amount of probability from pW on a vertex u can be written as follows:

$$\begin{aligned} pW(u) &= \frac{1}{2}p(u) + \frac{1}{2} \sum_{(v,u) \in E} \frac{p(v)}{d(v)} \\ &= \frac{1}{2} \sum_{(u,v) \in E} p(u,v) + \frac{1}{2} \sum_{(v,u) \in E} p(v,u) \\ &= \frac{1}{2}p(\text{in}(u)) + \frac{1}{2}p(\text{out}(u)). \end{aligned}$$

The amount of probability on a set S can then be written as

$$\begin{aligned} pW(S) &= \frac{1}{2}p(\text{in}(S)) + \frac{1}{2}p(\text{out}(S)) \\ &= \frac{1}{2}p(\text{in}(S) \cup \text{out}(S)) + \frac{1}{2}p(\text{in}(S) \cap \text{out}(S)). \end{aligned}$$

□

Now consider a PageRank vector defined by the usual equation,

$$\text{pr}_\alpha(s) = \alpha s + (1 - \alpha)\text{pr}_\alpha(s)W.$$

This equation relates $\text{pr}_\alpha(s)$ to $\text{pr}_\alpha(s)W$. By writing $[\text{pr}_\alpha(s)W](S)$ in terms of $\text{pr}_\alpha(s)$ using Lemma 4.2, we can relate $\text{pr}_\alpha(s)$ to itself. As a result, we can bound the height of the Lovasz-Simonovits curve for $\text{pr}_\alpha(s)$ in terms of other points on the same curve.

Lemma 4.3. *If $p = \text{pr}_\alpha(s)$ is a PageRank vector, then for any set S of vertices,*

$$p(S) = \alpha s(S) + (1 - \alpha) \left(\frac{1}{2} p(\text{in}(S) \cap \text{out}(S)) + \frac{1}{2} p(\text{in}(S) \cup \text{out}(S)) \right).$$

Furthermore, for each $j \in [1, n - 1]$,

$$\begin{aligned} p[\text{vol}(S_j^p)] &\leq \alpha s[\text{vol}(S_j^p)] \\ &\quad + (1 - \alpha) \left(\frac{1}{2} p[\text{vol}(S_j^p) - |\partial(S_j^p)|] + \frac{1}{2} p[\text{vol}(S_j^p) + |\partial(S_j^p)|] \right). \end{aligned}$$

Proof. Let $p = \text{pr}_\alpha(s - r)$ be an approximate PageRank vector. Since r is non-negative, we have the equation

$$\begin{aligned} p &= \alpha(s - r) + (1 - \alpha)pW \\ &\leq \alpha s + (1 - \alpha)pW. \end{aligned}$$

Applying Lemma 4.2, we obtain

$$\begin{aligned} p(S) &\leq \alpha s(S) + (1 - \alpha)pW(S) \\ &\leq \alpha s(S) + (1 - \alpha) \left(\frac{1}{2} p(\text{in}(S) \cup \text{out}(S)) + \frac{1}{2} p(\text{in}(S) \cap \text{out}(S)) \right). \end{aligned}$$

This proves the first part of the lemma. To prove the second part, recall that $p[\text{vol}(S_j^p)] = p(S_j^p)$ for any integer $j \in [0, n]$. It can be verified that the curve $p[x]$ gives an upper bound on the amount of probability moving over any set of directed edges A . Specifically, we have the bound $p(A) \leq p[|A|]$. Using these facts,

$$\begin{aligned} p[\text{vol}(S_j^p)] &= p(S_j^p) \\ &\leq \alpha s(S_j^p) + (1 - \alpha) \left(\frac{1}{2} p(\text{in}(S_j^p) \cup \text{out}(S_j^p)) \right. \\ &\quad \left. + \frac{1}{2} p(\text{in}(S_j^p) \cap \text{out}(S_j^p)) \right) \\ &\leq \alpha s[\text{vol}(S_j^p)] + (1 - \alpha) \left(\frac{1}{2} p[|\text{in}(S_j^p) \cup \text{out}(S_j^p)|] \right. \\ &\quad \left. + \frac{1}{2} p[|\text{in}(S_j^p) \cap \text{out}(S_j^p)|] \right). \end{aligned}$$

The lemma will follow by bounding the sizes of the sets in the inequality above. Notice that

$$|\text{in}(S_j^p) \cup \text{out}(S_j^p)| + |\text{in}(S_j^p) \cap \text{out}(S_j^p)| = 2\text{vol}(S_j^p),$$

and

$$|\text{in}(S_j^p) \cup \text{out}(S_j^p)| - |\text{in}(S_j^p) \cap \text{out}(S_j^p)| = 2|\partial(S_j^p)|.$$

This implies that

$$|\text{in}(S_j^p) \cup \text{out}(S_j^p)| = \text{vol}(S_j^p) + |\partial(S_j^p)|,$$

and

$$|\text{in}(S_j^p) \cap \text{out}(S_j^p)| = \text{vol}(S_j^p) - |\partial(S_j^p)|.$$

The result follows. \square

One consequence of the previous lemma is that a PageRank vector is more spread out than its starting vector in the sense that the curve $\text{pr}_\alpha(s)[x]$ lies below the curve $s[x]$. This monotonicity property is stated in the following lemma. We remark that this does not imply that $\text{pr}_\alpha(s)$ is less than or equal to s at any particular vertex.

Lemma 4.4. *For any starting vector s , and any $x \in [0, 2m]$,*

$$\text{pr}_\alpha(s)[x] \leq s[x].$$

Proof. If we let $p = \text{pr}_\alpha(s)$, Lemma 4.3 implies that for each $j \in [1, n-1]$,

$$\begin{aligned} p[\text{vol}(S_j^p)] &\leq \alpha s[\text{vol}(S_j^p)] + (1-\alpha) \left(\frac{1}{2} p[\text{vol}(S_j^p) - |\partial(S_j^p)|] \right. \\ &\quad \left. + \frac{1}{2} p[\text{vol}(S_j^p) + |\partial(S_j^p)|] \right) \\ &\leq \alpha s[\text{vol}(S_j^p)] + (1-\alpha)p[\text{vol}(S_j^p)], \end{aligned}$$

where the last line follows from the concavity of $p[k]$. This implies that $p[\text{vol}(S_j^p)] \leq s[\text{vol}(S_j^p)]$ for each $j \in [1, n-1]$. The same equation then holds for all $x \in [0, 2m]$, because $s[x]$ is concave and $p[x]$ is linear between these points. \square

Lemma 4.3 is stronger than the monotonicity property of Lemma 4.4, because it shows that for the PageRank vector $p = \text{pr}_\alpha(s)$, the height of the curve $p[x]$

at the point $\text{vol}(S_j^p)$ is at most the average of the heights of the same curve at the two points $\text{vol}(S_j^p) - |\partial(S_j^p)|$ and $\text{vol}(S_j^p) + |\partial(S_j^p)|$, plus a small term that depends on α . In the following theorem, we will use this averaging equation to derive a series of upper bounds on the curve $p[x]$. As a result, we can show that one of the following must be true: either $p(S) - \psi_V(S)$ is not very large for any set S of vertices in the graph, or else there is some sweep set S_j^p where $p(S_j^p) - \psi_V(S)$ is fairly large, and where S_j^p has small conductance.

Theorem 4.5. *Let $p = \text{pr}_\alpha(s)$ be a PageRank vector with $\|s_+\|_1 \leq 1$. Let ϕ and γ be any constants in $[0, 1]$. Either the bound*

$$p(S) - \psi_V(S) \leq \gamma + \alpha t + \sqrt{\min(\text{vol}(S), 2m - \text{vol}(S))} \left(1 - \frac{\phi^2}{8}\right)^t$$

holds for any set of vertices S and any integer t , or else there exists a sweep cut S_j^p , for some $j \in [1, |\text{Supp}(p)|]$, with the following properties:

1. $\Phi(S_j^p) < \phi$,
2. for some integer t ,

$$p(S_j^p) - \psi_V(S_j^p) > \gamma + \alpha t + \sqrt{\min(\text{vol}(S_j^p), 2m - \text{vol}(S_j^p))} \left(1 - \frac{\phi^2}{8}\right)^t.$$

Proof. Let $f_t(x) = \gamma + \alpha t + \sqrt{\min(x, 2m - x)} \left(1 - \frac{\phi^2}{8}\right)^t$, and consider the equation

$$p[x] - \frac{x}{2m} \leq f_t(x) \quad \text{for all } x \in [0, 2m]. \quad (4.1)$$

Assuming that there does not exist a sweep cut with both of the properties stated in the theorem, we will prove by induction that this equation holds for all $t \geq 0$. The theorem will follow.

For the base case $t = 0$, notice that for any value of x in the interval $[0, 2m]$,

$$p[x] - \frac{x}{2m} \leq \min(1, \min(x, 2m - x)) \leq \sqrt{\min(x, 2m - x)}.$$

Here, we have used the fact that $\|s_+\|_1 \leq 1$. This shows that Equation (4.1) holds at $t = 0$ for any choice of γ and ϕ .

Now assume for the sake of induction that Equation (4.1) holds for some specific t . To prove that Equation (4.1) holds for $t + 1$ at every point $x \in [0, 2m]$, it suffices to show that it holds for $t + 1$ at the points $x_j = \text{vol}(S_j^p)$ for each

$j \in [1, |\text{Supp}(p)]$. We already know that the equation holds trivially at $x_0 = 0$ and $x_n = 2m$. The result will follow because we will have shown that $p[x] - x/2m$ is piecewise linear between a set of points for which Equation (4.1) holds, and because $f_{t+1}(x)$ is concave.

Consider an index $j \in [1, |\text{Supp}(p)]$. We know that either Property 1 or Property 2 of the theorem does not hold for S_j . If Property 2 does not hold, this directly implies that Equation (4.1) holds with $t + 1$ and $x = x_j$. If Property 1 does not hold, then we have $\Phi(S_j^p) \geq \phi$, and we will apply Lemma 4.3. There are two cases to consider, depending on whether $x_j \leq m$ or $x_j \geq m$. We will carry out the proof assuming that $x_j \leq m$. The proof for the other case is similar.

$$\begin{aligned}
p[\text{vol}(S_j^p)] &\leq \alpha s[\text{vol}(S_j^p)] + (1 - \alpha) \left(\frac{1}{2} p[\text{vol}(S_j^p) - |\partial(S_j^p)|] \right. \\
&\quad \left. + \frac{1}{2} p[\text{vol}(S_j^p) + |\partial(S_j^p)|] \right) \\
&\leq \alpha + \left(\frac{1}{2} p[\text{vol}(S_j^p) - |\partial(S_j^p)|] + \frac{1}{2} p[\text{vol}(S_j^p) + |\partial(S_j^p)|] \right) \\
&= \alpha + \left(\frac{1}{2} p[x_j - \Phi(S_j^p)x_j] + \frac{1}{2} p[x_j + \Phi(S_j^p)x_j] \right) \\
&\leq \alpha + \left(\frac{1}{2} p[x_j - \phi x_j] + \frac{1}{2} p[x_j + \phi x_j] \right).
\end{aligned}$$

The last step above follows from the concavity of $p[x]$.

Using the induction hypothesis,

$$\begin{aligned}
p[x_j] &\leq \alpha + \frac{1}{2} \left(f_t(x_j - \phi x_j) + \frac{x_j - \phi x_j}{2m} + f_t(x_j + \phi x_j) + \frac{x_j + \phi x_j}{2m} \right) \\
&= \alpha + \frac{x_j}{2m} + \frac{1}{2} (f_t(x_j - \phi x_j) + f_t(x_j + \phi x_j)).
\end{aligned}$$

Therefore,

$$\begin{aligned}
p[x_j] - \frac{x_j}{2m} &\leq \alpha + \frac{1}{2} (f_t(x_j - \phi x_j) + f_t(x_j + \phi x_j)) \\
&= \gamma + \alpha + \alpha t + \frac{1}{2} \left(\sqrt{x_j - \phi x_j} + \sqrt{x_j + \phi x_j} \right) \left(1 - \frac{\phi^2}{8} \right)^t.
\end{aligned}$$

By examining the Taylor series of $\sqrt{1+\phi}$ at $\phi = 0$, we see that for any $x \geq 0$ and $\phi \in [0, 1]$,

$$\begin{aligned} & \frac{1}{2} \left(\sqrt{x - \phi x} + \sqrt{x + \phi x} \right) \\ & \leq \sqrt{x} \cdot \frac{1}{2} \left(\left(1 - \frac{1}{2}\phi - \frac{1}{8}\phi^2 - \frac{1}{16}\phi^3 - \dots \right) + \left(1 + \frac{1}{2}\phi - \frac{1}{8}\phi^2 + \frac{1}{16}\phi^3 - \dots \right) \right) \\ & \leq \sqrt{x} \left(1 - \frac{\phi^2}{8} \right). \end{aligned}$$

Therefore,

$$\begin{aligned} p[x_j] - \frac{x_j}{2m} & \leq \gamma + \alpha + \alpha t + \sqrt{x_j} \left(1 - \frac{\phi^2}{8} \right) \left(1 - \frac{\phi^2}{8} \right)^t \\ & = f_{t+1}(x_j). \end{aligned}$$

□

We will now derive Theorem 4.1.

Proof of Theorem 4.1. Let $\phi = \Phi(\text{pr}_\alpha(s))$. Theorem 4.5 implies that for any integer $t \geq 0$ and any set S ,

$$[\text{pr}_\alpha(s)](S) - \psi_V(S) \leq \alpha t + \sqrt{\min(\text{vol}(S), 2m - \text{vol}(S))} \left(1 - \frac{\phi^2}{8} \right)^t.$$

If we set

$$t = \left\lceil \frac{8}{\phi^2} \log(4\sqrt{\text{vol}(S)}/\delta) \right\rceil \leq \frac{9}{\phi^2} \log(4\sqrt{\text{vol}(S)}/\delta),$$

then we have

$$\sqrt{\min(\text{vol}(S), 2m - \text{vol}(S))} \left(1 - \frac{\phi^2}{8} \right)^t \leq \frac{\delta}{4}.$$

This gives the bound

$$[\text{pr}_\alpha(s)](S) - \psi_V(S) \leq \alpha \frac{9}{\phi^2} \log(4\sqrt{\text{vol}(S)}/\delta) + \frac{\delta}{4}.$$

On the other hand, we have assumed that $[\text{pr}_\alpha(s)](S) - \psi_V(S) > \delta$. Combining these upper and lower bounds yields the following inequality,

$$\frac{3\delta}{4} < \alpha \frac{9}{\phi^2} \log(4\sqrt{\text{vol}(S)}/\delta).$$

The result follows by solving for ϕ .

□

5. Cuts from PageRank Vectors

Consider the following procedure: pick a starting vertex v and a value of α , compute an ϵ -approximate PageRank vector for $\text{pr}_\alpha(s)$, and perform a sweep over the resulting approximation. In this section, we show that for any set C of conductance $O(\alpha)$ and for many of the vertices v within C , this procedure finds a set with conductance $O(\sqrt{\alpha \log(\text{vol}(C))})$. To prove this, we identify a set of vertices v within C for which we can give a lower bound on the amount of probability from $\text{pr}_\alpha(\chi_v)$ on the set C , and then we apply the results from the previous section.

Theorem 5.1. *For any set C and any constant α in $(0, 1]$, there is a subset $C_\alpha \subseteq C$ with volume $\text{vol}(C_\alpha) \geq \text{vol}(C)/2$ such that for any vertex $v \in C_\alpha$, the PageRank vector $\text{pr}_\alpha(\chi_v)$ satisfies*

$$[\text{pr}_\alpha(\chi_v)](C) \geq 1 - \frac{\Phi(C)}{\alpha}.$$

Proof. Theorem 5.1 states that a set C with small conductance contains a significant amount of probability from $\text{pr}_\alpha(\chi_v)$ for many of the vertices v in C . We first show that this holds for an average of the vertices in C , by showing that the PageRank vector $\text{pr}_\alpha(\psi_C)$ satisfies

$$[\text{pr}_\alpha(\psi_C)](\bar{C}) \leq \Phi(C) \frac{1 - \alpha}{2\alpha}. \quad (5.1)$$

We then observe that if we sample a vertex from the distribution ψ_C , at least half of the time $\text{pr}_\alpha(\chi_v)$ is less than twice its expectation, which is $\text{pr}_\alpha(\psi_C)$.

We will prove Equation (5.1) by examining the movement of probability during the single step from $\text{pr}_\alpha(\psi_C)$ to $\text{pr}_\alpha(\psi_C)W$. The amount of probability that moves from C to \bar{C} in the step from $\text{pr}_\alpha(\psi_C)$ to $\text{pr}_\alpha(\psi_C)W$ is bounded by $\frac{1}{2}\text{pr}_\alpha(\psi_C)[|\partial(C)|]$, so

$$[\text{pr}_\alpha(\psi_C)W](\bar{C}) \leq [\text{pr}_\alpha(\psi_C)](\bar{C}) + \frac{1}{2}\text{pr}_\alpha(\psi_C)[|\partial(C)|].$$

We combine this observation with the PageRank equation to obtain the following:

$$\begin{aligned} [\text{pr}_\alpha(\psi_C)](\bar{C}) &= [\alpha\psi_C + (1 - \alpha)\text{pr}_\alpha(\psi_C)W](\bar{C}) \\ &= (1 - \alpha)[\text{pr}_\alpha(\psi_C)W](\bar{C}) \\ &\leq (1 - \alpha)[\text{pr}_\alpha(\psi_C)](\bar{C}) + \frac{1 - \alpha}{2}\text{pr}_\alpha(\psi_C)[|\partial(C)|]. \end{aligned}$$

This implies that

$$[\text{pr}_\alpha(\psi_C)](\bar{C}) \leq \frac{1-\alpha}{2\alpha} \text{pr}_\alpha(\psi_C)[|\partial(C)|].$$

This equation bounds the amount of probability outside of C in terms of $\frac{1}{2}\text{pr}_\alpha(\psi_C)[|\partial(C)|]$, which is an upper bound on the amount of probability that leaves C at each step. This quantity can be bounded in terms of the conductance of C . Using the monotonicity property from Lemma 4.4,

$$\begin{aligned} \text{pr}_\alpha(\psi_C)[|\partial(C)|] &\leq \psi_C[|\partial(C)|] \\ &= \frac{|\partial(C)|}{\text{vol}(C)} \\ &= \Phi(C). \end{aligned}$$

This implies that

$$[\text{pr}_\alpha(\psi_C)](\bar{C}) \leq \frac{1-\alpha}{2\alpha} \Phi(C).$$

To complete the proof, let C_α be the set of vertices v in C satisfying

$$\text{pr}_\alpha(\chi_v)(\bar{C}) \leq \frac{\Phi(C)}{\alpha}.$$

Let v be a vertex chosen randomly from the distribution ψ_C , and define the random variable $X = \text{pr}_\alpha(\chi_v)(\bar{C})$. The linearity property of PageRank vectors from Proposition 2.2 implies the following bound on the expectation of X :

$$\mathbb{E}[X] = \text{pr}_\alpha(\psi_C)(\bar{C}) \leq \frac{1-\alpha}{2\alpha} \Phi(C) \leq \frac{\Phi(C)}{2\alpha}.$$

Applying Markov's inequality yields

$$\Pr[v \notin C_\alpha] \leq \Pr[X > 2\mathbb{E}[X]] \leq \frac{1}{2}.$$

Since $\Pr[v \in C_\alpha] \geq \frac{1}{2}$, the volume of C_α is at least $\frac{1}{2}\text{vol}(C)$. □

We can give a similar lower bound on the amount of probability within C from an ϵ -approximate PageRank vector. We use the following lemma to bound the amount of probability that is lost in the approximation.

Lemma 5.2. *For any ϵ -approximate PageRank vector $\text{pr}_\alpha(s-r)$ and any set S of vertices,*

$$[\text{pr}_\alpha(s)](S) \geq [\text{pr}_\alpha(s-r)](S) \geq [\text{pr}_\alpha(s)](S) - \epsilon \text{vol}(S).$$

Proof. Since the vector r is nonnegative,

$$\text{pr}_\alpha(s - r) = \text{pr}_\alpha(s) - \text{pr}_\alpha(r) \leq \text{pr}_\alpha(s).$$

To prove the other half of the lemma, we use the monotonicity property from Lemma 4.4 to bound the difference between $\text{pr}_\alpha(s)$ and an ϵ -approximate PageRank vector $\text{pr}_\alpha(s - r)$. For any set S , we have

$$[\text{pr}_\alpha(r)](S) \leq \text{pr}_\alpha(r) [\text{vol}(S)] \leq r [\text{vol}(S)],$$

and so

$$\begin{aligned} [\text{pr}_\alpha(s - r)](S) &= [\text{pr}_\alpha(s)](S) - [\text{pr}_\alpha(r)](S) \\ &\geq [\text{pr}_\alpha(s)](S) - r [\text{vol}(S)] \\ &\geq [\text{pr}_\alpha(s)](S) - \epsilon \text{vol}(S). \end{aligned}$$

□

We now know that if v is a vertex in C_α and if $\text{pr}_\alpha(\chi_v - r)$ is an ϵ -approximate PageRank vector, then

$$[\text{pr}_\alpha(\chi_v - r)](C) \geq 1 - \frac{\Phi(C)}{\alpha} - \epsilon \text{vol}(C).$$

If both $\Phi(C)/\alpha$ and ϵ are small, there is still a significant amount of probability on the set C , so we can apply the mixing result from Theorem 4.1 to show that a sweep over $\text{pr}_\alpha(\chi_v - r)$ finds a cut with small conductance.

Theorem 5.3. *Let α be a constant in $(0, 1]$, and let C be a set satisfying*

1. $\Phi(C) \leq \alpha/10$,
2. $\text{vol}(C) \leq \frac{2}{3} \text{vol}(G)$.

If $\tilde{p} = \text{pr}_\alpha(\chi_v - r)$ is an ϵ -approximate PageRank vector where $v \in C_\alpha$ and $\epsilon \leq \frac{1}{10 \text{vol}(C)}$, then a sweep over \tilde{p} produces a cut with conductance $\Phi(\tilde{p}) = O(\sqrt{\alpha \log(\text{vol}(C))})$.

Proof. Let $\tilde{p} = \text{pr}_\alpha(\chi_v - r)$ be an ϵ -approximate PageRank vector for $\text{pr}_\alpha(\chi_v)$ satisfying the assumptions of the theorem. Combining Theorem 5.1 with Lemma 5.2 gives a lower bound on $\tilde{p}(C)$,

$$\tilde{p}(C) \geq 1 - \frac{\Phi(C)}{\alpha} - \epsilon \text{vol}(C).$$

Since $\Phi(C)/\alpha \leq 1/10$ and $\epsilon \leq 1/(10\text{vol}(C))$, we have $\tilde{p}(C) \geq 4/5$, which implies

$$\tilde{p}(C) - \psi_V(C) \geq \frac{4}{5} - \frac{2}{3} = \frac{2}{15}.$$

Theorem 4.1 then implies

$$\Phi(\tilde{p}) < \sqrt{90\alpha \log(30\sqrt{\text{vol}(C)})}.$$

□

As a corollary, there is some starting vertex v and value of α for which a sweep over $\text{pr}_\alpha(\chi_v)$ finds a cut with conductance near the minimum conductance in the graph.

Corollary 5.4. *Let Φ_G be the minimum conductance of any set of vertices in the graph, and let C^{opt} be a set achieving this minimum. If \tilde{p} is an ϵ -approximate PageRank vector for $\text{pr}_\alpha(\chi_v)$, where $\alpha = 10\Phi_G$, $v \in C_\alpha^{\text{opt}}$, and $\epsilon \leq \frac{1}{10\text{vol}(C^{\text{opt}})}$, then*

$$\Phi(\tilde{p}) = O(\sqrt{\Phi_G \log(\text{vol}(C^{\text{opt}})})).$$

Corollary 5.4 follows from Theorem 5.3 by setting $C = C^{\text{opt}}$.

6. Finding Small Cuts in Nearly Linear Time

In the previous section, we showed that to find a cut within a set C , it suffices to compute an ϵ -approximate PageRank vector with ϵ roughly $1/\text{vol}(C)$. This requires time proportional to $\text{vol}(C)$, but the resulting cut S may have much smaller volume. This leaves us with a problem similar to the one facing recursive spectral partitioning: we do not want to spend a large amount of time to find a cut whose volume is small.

In this section, we extend our local partitioning techniques to find a cut with small conductance in time proportional to the volume of the smaller side of the cut found. Essentially, we consider what would happen if we were to compute an approximate PageRank vector with different levels of error. We show that if a given level of error yields a cut whose volume is too small, we could have found a cut with similar volume more quickly by using a larger amount of error.

The result is an algorithm called PageRank-Nibble, which takes a constant $\phi \in (0, 1]$ and a scale $b \in [1, \log m]$ as part of its input and attempts to find a cut with conductance ϕ and volume at least 2^{b-1} . We prove that PageRank-Nibble finds a set with these properties for at least one choice of b in the range

Algorithm 2. (PageRank-Nibble (v, ϕ, b) .)

Input: a vertex v , a constant $\phi \in (0, 1]$, and an integer $b \in [1, B]$, where $B = \lceil \log m \rceil$.

1. Let $\alpha = \frac{\phi^2}{225 \log(100\sqrt{m})}$.
 2. Compute an ϵ -approximate PageRank vector $p = \text{pr}_\alpha(\chi_v - r)$, with $\epsilon \leq (2^b \cdot 48 \lceil \log m \rceil)^{-1}$.
 3. For each $j \in [1, |\text{Supp}(p)|]$, check whether S_j^p obeys the following conditions:
 - conductance: $\Phi(S_j^p) < \phi$,
 - volume: $2^{b-1} < \text{vol}(S_j^p) < \frac{2}{3} \text{vol}(G)$,
 - probability: $p[2^b] - p[2^{b-1}] > \frac{1}{48 \lceil \log m \rceil}$.
 4. If some set S_j^p satisfies all of these conditions, return S_j^p . Otherwise, return nothing.
-

$[1, \lceil \log m \rceil]$ whenever v is a good starting vertex for a set C with sufficiently small conductance. In addition, we show that the resulting set has a large intersection with the set C .

Theorem 6.1. *PageRank-Nibble* (v, ϕ, b) (Algorithm 2) can be implemented with running time $O(2^b \frac{\log^2 m}{\phi^2})$.

Proof. An ϵ -approximate PageRank vector p with $\epsilon \leq (2^b \cdot 48 \lceil \log m \rceil)^{-1}$ can be computed in time $O(2^b \frac{\log m}{\alpha})$ using ApproximatePR. By Theorem 3.2, the support of this vector has volume $O(2^b \log m)$, and the number of vertices in the support is $N_p = O(2^b \log m)$. It is possible to check each of the conditions in step 3 of PageRank-Nibble, for every set S_j^p with $j \in [1, N_p]$, in the amount of time required to sort and perform a sweep, which is

$$O(\text{vol}(\text{Supp}(p)) + N_p \log N_p) = O(2^b \log^2 m).$$

Since we have set $\alpha = \Omega(\phi^2 / \log m)$, the running time of PageRank-Nibble is

$$O(2^b \frac{\log m}{\alpha} + 2^b \log^2 m) = O(2^b \frac{\log^2 m}{\phi^2}).$$

□

Theorem 6.2. *Let C be a set such that*

$$\text{vol}(C) \leq \frac{1}{2} \text{vol}(G)$$

and $\Phi(C) \leq \phi^2 / (22500 \log^2 100m)$, and let v be a vertex in C_α for the value of α used in PageRank-Nibble, which is $\phi^2 / (225 \log(100\sqrt{m}))$. Then, there is some integer $b \in [1, \lceil \log m \rceil]$ for which PageRank-Nibble(v, ϕ, b) finds a set S meeting all of its criteria. Any such set has the following properties:

- *conductance:* $\Phi(S) < \phi$,
- *volume:* $2^{b-1} < \text{vol}(S) < \frac{2}{3} \text{vol}(G)$,
- *intersection:* $\text{vol}(S \cap C) > 2^{b-2}$.

Proof. Consider the PageRank vector $\text{pr}_\alpha(\chi_v)$. We have assumed that v is in C_α and have set $\Phi(C)$ to ensure that $\Phi(C) \leq \alpha / (100 \lceil \log m \rceil)$. This implies that

$$\begin{aligned} \text{pr}_\alpha(\chi_v) [\text{vol}(C)] - \psi_V(C) &\geq \left(1 - \frac{\phi(C)}{\alpha}\right) - \frac{1}{2} \\ &\geq \frac{1}{2} - \frac{1}{100}. \end{aligned}$$

We have set α so that if $t_0 = \lceil \frac{8}{\phi^2} \log(100\sqrt{m}) \rceil$, then $\alpha t_0 \leq 1/25$. With this choice of t_0 , we have

$$\alpha t_0 + \sqrt{\text{vol}(C)} \left(1 - \frac{\phi^2}{8}\right)^{t_0} < \frac{1}{25} + \frac{1}{100}.$$

Since $\frac{1}{2} - \frac{1}{100} > \frac{5}{12} + \frac{1}{25} + \frac{1}{100}$, the following equation holds:

$$\text{pr}_\alpha(\chi_v) [\text{vol}(C)] - \frac{\text{vol}(C)}{2m} > \frac{5}{12} + \alpha t_0 + \sqrt{\text{vol}(C)} \left(1 - \frac{\phi^2}{8}\right)^{t_0}. \quad (6.1)$$

Let $B = \lceil \log m \rceil$. For each integer b in $[1, B]$, let $\gamma_b = \frac{5}{12} \left(\frac{9}{10} + \frac{1}{10} \frac{b}{B}\right)$. We will consider the following equation:

$$\text{pr}_\alpha(\chi_v) [x] - \frac{x}{2m} > \gamma_b + \alpha t + \sqrt{x} \left(1 - \frac{\phi^2}{8}\right)^t. \quad (6.2)$$

We have already shown that this holds with $b = B$, $x = m$, and $t = t_0$, in Equation (6.1). Let b_0 be the smallest value of b for which this equation holds for some $x_0 \leq 2^b$ and for some value of t . We will show that PageRank-Nibble successfully returns a cut when it is run with $b = b_0$.

When PageRank-Nibble is run with $b = b_0$, it computes an ϵ -approximate PageRank vector $\text{pr}_\alpha(\chi_v - r)$ with $\epsilon \leq (2^{b_0} 48B)^{-1}$. With this amount of error, we have

$$\begin{aligned} \text{pr}_\alpha(\chi_v - r)[x_0] &\geq \text{pr}_\alpha(\chi_v)[x_0] - \left(\frac{2^{-b_0}}{48B}\right) x_0 \\ &\geq \text{pr}_\alpha(\chi_v)[x_0] - \frac{1}{48B} \\ &\geq \text{pr}_\alpha(\chi_v)[x_0] - (\gamma_{b_0} - \gamma_{b_0-1}) + \frac{1}{48B}, \end{aligned}$$

where the last line follows because $\gamma_{b_0} - \gamma_{b_0-1} \leq \frac{1}{24B}$. Since Equation (6.2) holds for b_0 and x_0 , we have for some integer $t \geq 0$,

$$\begin{aligned} \text{pr}_\alpha(\chi_v - r)[x_0] - \frac{x_0}{2m} &> \left(\gamma_{b_0} + \alpha t + \sqrt{x_0} \left(1 - \frac{\phi^2}{8}\right)^t \right) \\ &\quad - (\gamma_{b_0} - \gamma_{b_0-1}) + \frac{1}{48B} \\ &> \left(\gamma_{b_0-1} + \frac{1}{48B} \right) + \alpha t + \sqrt{x_0} \left(1 - \frac{\phi^2}{8}\right)^t. \end{aligned}$$

Theorem 4.5 then shows that there exists a sweep cut S_j , with $S_j = S_j^{\text{pr}_\alpha(\chi_v - r)}$ for some value of j in the range $[1, |\text{Supp}(\text{pr}_\alpha(\chi_v - r))|]$, such that $\Phi(S_j) \leq \phi$, and such that the following lower bound holds for some integer t' :

$$\text{pr}_\alpha(\chi_v - r)(S_j) - \frac{\text{vol}(S_j)}{2m} > \left(\gamma_{b_0-1} + \frac{1}{48B} \right) + \alpha t' + \sqrt{\overline{\text{vol}}(S_j)} \left(1 - \frac{\phi^2}{8}\right)^{t'}, \quad (6.3)$$

where $\overline{\text{vol}}(S_j) = \min(\text{vol}(S_j), 2m - \text{vol}(S_j))$. We will show that this cut S_j meets all the requirements of PageRank-Nibble, which will prove that the algorithm outputs some cut when run with $b = b_0$.

First, assume for the sake of contradiction that $\text{vol}(S_j) \leq 2^{b_0-1}$. Since Equation (6.2) cannot hold with $b = b_0 - 1$ and $x \leq 2^{b_0-1}$, this implies that for any integer $t \geq 0$,

$$\begin{aligned} \text{pr}_\alpha(\chi_v - r)(S_j) - \frac{\text{vol}(S_j)}{2m} &= \text{pr}_\alpha(\chi_v - r)[\text{vol}(S_j)] - \frac{\text{vol}(S_j)}{2m} \\ &\leq \text{pr}_\alpha(\chi_v)[\text{vol}(S_j)] - \frac{\text{vol}(S_j)}{2m} \\ &\leq \gamma_{b_0-1} + \alpha t + \sqrt{\text{vol}(S_j)} \left(1 - \frac{\phi^2}{8}\right)^t. \end{aligned}$$

Since $\overline{\text{vol}}(S_j) = \text{vol}(S_j)$ when $x \leq 2^{b_0-1}$, this contradicts the lower bound from Equation (6.3). Therefore, it must be true that $\text{vol}(S_j) > 2^{b_0-1}$.

It must also be true that $\text{vol}(S_j) < \frac{2}{3}\text{vol}(G)$. Otherwise, the lower bound from Equation (6.3) would imply that for some integer $t' \geq 0$,

$$\begin{aligned} \text{pr}_\alpha(\chi_v - r)(S_j) &> \frac{\text{vol}(S_j)}{2m} + \gamma_{b_0-1} + \alpha t' + \sqrt{\text{vol}(S_j)} \left(1 - \frac{\phi^2}{8}\right)^{t'} \\ &> \frac{2}{3} + \gamma_{b_0-1} \\ &\geq \frac{2}{3} + \frac{9}{10} \frac{5}{12}. \end{aligned}$$

This implies $\text{pr}_\alpha(\chi_v - r)(S_j) > 1$, which is a contradiction.

We will now prove that there is a significant difference in probability between $\text{pr}_\alpha(\chi_v - r)[2^{b_0}]$ and $\text{pr}_\alpha(\chi_v - r)[2^{b_0-1}]$. Since Equation (6.3) does not hold with $b = b_0 - 1$ and $x = 2^{b_0-1}$, we know that for every integer $t \geq 0$,

$$\text{pr}_\alpha(\chi_v - r)[2^{b_0-1}] - \frac{x_0}{2m} \leq \gamma_{b_0-1} + \alpha t + \sqrt{2^{b_0-1}} \left(1 - \frac{\phi^2}{8}\right)^t. \quad (6.4)$$

We also know that for some integer t' ,

$$\text{pr}_\alpha(\chi_v - r)[x_0] - \frac{x_0}{2m} > \left(\gamma_{b_0-1} + \frac{1}{48B}\right) + \alpha t' + \sqrt{x_0} \left(1 - \frac{\phi^2}{8}\right)^{t'}. \quad (6.5)$$

By plugging t' into Equation (6.4), we obtain the following inequality:

$$\begin{aligned} \text{pr}_\alpha(\chi_v - r)[2^{b_0}] - \text{pr}_\alpha(\chi_v - r)[2^{b_0-1}] &\geq \text{pr}_\alpha(\chi_v - r)[x_0] \\ &\quad - \text{pr}_\alpha(\chi_v - r)[2^{b_0-1}] \\ &> \frac{1}{48B}. \end{aligned}$$

We have shown that S_j meets all the requirements of PageRank-Nibble, which proves that the algorithm outputs some cut when run with $b = b_0$. We now prove a lower bound on $\text{vol}(S \cap C)$, which holds for any cut S output by PageRank-Nibble, regardless of whether the algorithm was run with $b = b_0$ or with some other value of b . Let $p'[x] = p[x] - p[x-1]$. Since $p'[x]$ is a decreasing function of x ,

$$\begin{aligned} p'[2^{b-1}] &\geq \frac{p[2^b] - p[2^{b-1}]}{2^b - 2^{b-1}} \\ &> \frac{1}{2^{(b-1)}48B}. \end{aligned}$$

It is not hard to see that combining this lower bound on $p' [2^{b-1}]$ with the upper bound $p(\bar{C}) \leq \frac{\Phi(C)}{\alpha}$ gives the following bound on the volume of the intersection:

$$\begin{aligned} \text{vol}(S_j \cap C) &\geq 2^{b-1} - \frac{p(\bar{C})}{p' [2^{b-1}]} \\ &> 2^{b-1} - 2^{b-1} \left(48B \frac{\Phi(C)}{\alpha} \right). \end{aligned}$$

Since we have assumed that $\frac{\Phi(C)}{\alpha} \leq \frac{1}{100B}$, we have

$$\text{vol}(S \cap C) > 2^{b-1} - 2^{b-2} = 2^{b-2}.$$

□

7. Local Graph Partitioning

PageRank-Nibble improves the running time and approximation ratio of the Nibble algorithm of Spielman and Teng [Spielman and Teng 04]. In their paper, Nibble was called repeatedly with randomly chosen starting vertices and scales to create an algorithm called Partition, which finds a cut with small conductance and approximately optimal volume. Partition was applied recursively to create algorithms for multiway partitioning, graph sparsification, and solving diagonally dominant linear systems.

The PageRank-Partition algorithm can be created by calling PageRank-Nibble instead of Nibble. The algorithm takes as input a parameter ϕ and a graph, and it has expected running time $O(m \log^4 m / \phi^2)$. If there exists a set C with $\Phi(C) = O(\phi^2 / \log^2 m)$, then with high probability PageRank-Partition finds a set S such that $\text{vol}(S) \geq \text{vol}(C)/2$ and $\Phi(S) \leq \phi$.

In Table 1, we compare our local partitioning algorithms with the existing ones. The running times are stated in terms of ϕ , which is the conductance of the cut returned by the algorithm. The approximation ratios are described by stating what $\Phi(C)$ must be to guarantee that the algorithm will find a cut of conductance ϕ with high probability.

	Running time	Approximation
Nibble	$2^b \log^4 m / \phi^5$	$\phi^3 / \log^2 m$
PR-Nibble	$2^b \log^2 m / \phi^2$	$\phi^2 / \log^2 m$
Partition	$m \log^6 m / \phi^5$	$\phi^3 / \log^2 m$
PR-Partition	$m \log^4 m / \phi^2$	$\phi^2 / \log^2 m$

Table 1. Comparison of local partitioning algorithms.




Figure 1. Two orderings of the adjacency matrix of a graph derived from the Internet Movie Database, created by applying spectral partitioning (left) and PageRank-Nibble (right). Each dot represents an edge in the graph, which represents an appearance of an actress in a movie. The shade of the dot is determined by the country in which the movie was produced.

Finding balanced cuts in nearly linear time with PageRank-Partition is one important application of our local partitioning techniques. Recently, Khandekar, Rao, and Vazirani, introduced an algorithm that produces balanced cuts quickly using a different method [Khandekar et al. 06]. Their algorithm produces an $O(\log^2 n)$ approximation for the balanced cut problem using $O(\log^4 n)$ single commodity flow computations.

Figure 1 depicts two orderings of the adjacency matrix of a bipartite graph derived from the Internet Movie Database.¹ The graph contains 198,430 nodes, each representing either an actress or a movie, and 1,133,512 edges, with an edge appearing between a movie and an actress if and only if the actress appears in the movie. The ordering on the left was created by applying spectral partitioning recursively. The ordering on the right was created by applying PageRank-Nibble many times with random starting vertices and scales to obtain a large collection of cuts. Each vertex was then assigned the value of the smallest conductance cut from the collection of which it was a member, and the vertices were ordered from left to right in increasing order of this value. Both orderings make the adjacency matrix of the movie graph approximately block-diagonal, with blocks corresponding to different countries.

¹<http://www.imdb.com/>

8. Appendix

The following proposition shows that the definition of PageRank used in this paper, which uses the lazy random walk matrix $W = \frac{1}{2}(I + D^{-1}A)$, is equivalent to the standard definition, which uses the random walk matrix $M = D^{-1}A$.

Proposition 8.1. *Let $p = \text{pr}_\alpha(s)$ be the unique solution to the equation*

$$p = \alpha s + (1 - \alpha)pW,$$

where $W = \frac{1}{2}(I + D^{-1}A)$. Then, p is also the unique solution of the traditional PageRank equation

$$p = \alpha' s + (1 - \alpha')pM,$$

where $\alpha' = 2\alpha/(1 + \alpha)$ and $M = D^{-1}A$.

Proof Proof. The proof is just algebra.

$$\begin{aligned} \text{pr}_\alpha(s) &= \alpha s + (1 - \alpha)\text{pr}_\alpha(s)W \\ &= \alpha s + \left(\frac{1 - \alpha}{2}\right) \text{pr}_\alpha(s) + \left(\frac{1 - \alpha}{2}\right) \text{pr}_\alpha(s)(D^{-1}A). \end{aligned}$$

This implies that

$$\left(\frac{1 + \alpha}{2}\right) \text{pr}_\alpha(s) = \alpha s + \left(\frac{1 - \alpha}{2}\right) \text{pr}_\alpha(s)(D^{-1}A),$$

and so

$$\text{pr}_\alpha(s) = \left(\frac{2\alpha}{1 + \alpha}\right) s + \left(1 - \frac{2\alpha}{1 + \alpha}\right) \text{pr}_\alpha(s)(D^{-1}A).$$

The result follows. □

Proof of Proposition 2.1. The solutions of the PageRank equation $p = \alpha s + (1 - \alpha)pW$ are the solutions of the linear system $p(I - (1 - \alpha)W) = \alpha s$. The matrix $(I - (1 - \alpha)W)$ is nonsingular, since it is strictly diagonally dominant, so this equation has a unique solution. □

Proof of Proposition 2.2. The sum that defines R_α in Equation (2.2) is absolutely convergent for $\alpha \in (0, 1]$, and the following computation shows that sR_α obeys

the steady state equation for $\text{pr}_\alpha(s)$:

$$\begin{aligned}\alpha s + (1 - \alpha)sR_\alpha W &= \alpha s + (1 - \alpha)s \left(\alpha I + \alpha \sum_{t=1}^{\infty} (1 - \alpha)^t W^t \right) W \\ &= \alpha s + s \left(\alpha \sum_{t=1}^{\infty} (1 - \alpha)^t W^t \right) \\ &= sR_\alpha.\end{aligned}$$

Since the solution to the this equation is unique by Proposition 2.1, it follows that $\text{pr}_\alpha(s) = sR_\alpha$. \square

Proof of Equation 3.1. The following sequence of equations shows that $\text{pr}_\alpha(s)W$ obeys the PageRank equation for $\text{pr}_\alpha(sW)$. This equation has a unique solution, and so $\text{pr}_\alpha(s)W = \text{pr}_\alpha(sW)$:

$$\begin{aligned}\text{pr}_\alpha(s) &= \alpha s + (1 - \alpha)\text{pr}_\alpha(s)W \\ \text{pr}_\alpha(s)W &= \alpha sW + (1 - \alpha)\text{pr}_\alpha(s)W^2 \\ (\text{pr}_\alpha(s)W) &= \alpha(sW) + (1 - \alpha)(\text{pr}_\alpha(s)W)W.\end{aligned}$$

\square

Acknowledgements. An extended abstract of this paper appeared in *Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS 2006)*.

References

- [Berkhin 06] Pavel Berkhin. “Bookmark-Coloring Algorithm for Personalized PageRank Computing.” *Internet Mathematics* 3:1 (2006), 41–62.
- [Borgs et al. 04] Christian Borgs, Jennifer T. Chayes, Mohammad Mahdian, and Amin Saberi. “Exploring the Community Structure of Newsgroups.” In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 783–787. New York: ACM Press, 2004.
- [Brin and Page 98] Sergey Brin and Lawrence Page. “The Anatomy of a Large-Scale Hypertextual Web Search Engine.” *Computer Networks and ISDN Systems* 30:1–7 (1998), 107–117.
- [Fogaras and Racz 04] D. Fogaras and B. Racz. “Towards Scaling Fully Personalized PageRank.” In *Algorithms and Models for the Web-Graph: Third International Workshop, WAW 2004, Rome, Italy, October 16, 2004, Proceedings*, Lecture Notes in Computer Science 3243, pp. 105–117. Berlin: Springer, 2004.

- [Haveliwala 03] Taher H. Haveliwala. “Topic-Sensitive PageRank: A Context-Sensitive Ranking Algorithm for Web Search.” *IEEE Trans. Knowl. Data Eng.* 15:4 (2003), 784–796.
- [Jeh and Widom 03] Glen Jeh and Jennifer Widom. “Scaling Personalized Web Search.” In *Proceedings of the 12th International Conference on World Wide Web*, pp. 271–279. New York: ACM Press, 2003.
- [Kannan et al. 04] Ravi Kannan, Santosh Vempala, and Adrian Vetta. “On Clusterings: Good, Bad and Spectral.” *J. ACM* 51:3 (2004), 497–515.
- [Khandekar et al. 06] Rohit Khandekar, Satish Rao, and Umesh Vazirani. “Graph Partitioning Using Single Commodity Flows.” In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, pp. 385–390. New York: ACM Press, 2006.
- [Leighton and Rao 88] Frank Thomson Leighton and Satish Rao. “An Approximate Max-Flow Min-Cut Theorem for Uniform Multicommodity Flow Problems with Applications to Approximation Algorithms.” In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pp. 422–431. Los Alamitos, CA: IEEE Press, 1988.
- [Lovász and Simonovits 90] László Lovász and Miklós Simonovits. “The Mixing Rate of Markov Chains, an Isoperimetric Inequality, and Computing the Volume.” In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*, pp. 346–354. Los Alamitos, CA: IEEE Press, 1990.
- [Lovász and Simonovits 93] László Lovász and Miklós Simonovits. “Random Walks in a Convex Body and an Improved Volume Algorithm.” *Random Struct. Algorithms* 4:4 (1993), 359–412.
- [Mihail 89] M. Mihail. “Conductance and Convergence of Markov Chains—A Combinatorial Treatment of Expanders.” In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, pp. 526–531. Los Alamitos, CA: IEEE Press, 1989.
- [Page et al. 98] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. “The PageRank Citation Ranking: Bringing Order to the Web.” Technical report, Stanford Digital Library Technologies Project, 1998.
- [Simon and Teng 97] Horst D. Simon and Shang-Hua Teng. “How Good is Recursive Bisection?” *SIAM Journal on Scientific Computing* 18:5 (1997), 1436–1445.
- [Spielman and Teng 96] Daniel A. Spielman and Shang-Hua Teng. “Spectral Partitioning Works: Planar Graphs and Finite Element Meshes.” In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pp. 96–105. Los Alamitos, CA: IEEE Press, 1996.
- [Spielman and Teng 04] Daniel A. Spielman and Shang-Hua Teng. “Nearly-Linear Time Algorithms for Graph Partitioning, Graph Sparsification, and Solving Linear Systems.” In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, pp. 81–90. New York: ACM Press, 2004.

Reid Andersen, Microsoft Live Labs, One Microsoft Way, Redmond, WA 98052
(reidan@microsoft.com)

Fan Chung, University of California, San Diego, Department of Mathematics, 9500
Gilman Drive, 0012, La Jolla, CA 92093-0112 (fan@ucsd.edu)

Kevin Lang, Yahoo! Research, 2821 Mission College Blvd., Santa Clara, CA 95054
(langk@yahoo-inc.com)

Received January 8 2007; accepted October 15, 2007.