

Dirichlet PageRank and Ranking Algorithms Based on Trust and Distrust

Fan Chung, Alexander Tsiatas, and Wensong Xu

Abstract. Motivated by numerous models of representing trust and distrust within a network ranking system, we examine a quantitative vertex ranking with consideration of the influence of a subset of nodes. We propose and analyze a general ranking metric, called *Dirichlet PageRank*, which gives a ranking of vertices in a subset S of nodes subject to some specified conditions on the vertex boundary of S . In addition to the usual Dirichlet boundary condition (which disregards the influence of nodes outside of S), we consider general boundary conditions allowing the presence of negative (distrustful) nodes or edges. We give an efficient approximation algorithm for computing Dirichlet PageRank vectors. Furthermore, we give several algorithms for solving various trust-based ranking problems using Dirichlet PageRank with general boundary conditions.

1. Introduction

PageRank has proven to be a useful tool for vertex ranking in many contexts, but some refinements are needed to address many increasingly complex but crucial problems. For example, PageRank is susceptible to manipulation by link

spammers, and it treats all links between nodes as positive votes for importance even if some links are meant to show distrust.

To illustrate the need for incorporating several different types of “trust” and “distrust,” we consider the following four examples:

Problem 1.1. Suppose a small community holds an election. A community can be represented by a subgraph in a social network, and only edges incident to nodes in the subgraph can be used to determine the ranking of the nodes. The original interpretation of PageRank treats each edge as a vote for determining the “importance” of the nodes. A local community’s election should not be influenced by interests outside of the community; one way to deal with this is to declare the influence of all outside nodes to be zero. This is the *Dirichlet boundary condition* that we will discuss in this paper.

Problem 1.2. In the web graph, there are many nodes whose importance is not properly reflected by the link structure of the graph. For example, the websites of some governmental agencies are known to have high impact and authority, but they may not be highly connected to other websites. With prior knowledge of the network, it is often desirable to be able to effectively adjust the ranking of exceptional webpages.

Problem 1.3. Another factor that many ranking models should address is the notion of “distrust.” Distrust can appear in many different ways; for example, if several vertices are known to be spammers, their neighbors are likely to be spammers as well. It is desirable to be able to quantify distrust, and such quantification can then be used for protection as well as for penalizing spammers. In some cases, distrust between vertices can be built into the graph as negatively weighted edges, presenting computational challenges that the usual PageRank definition does not address. As we will see in later sections, we can use negative links to build a new network with boundary conditions chosen to appropriately propagate distrust.

Problem 1.4. Another vertex-ranking problem arises from a distinction between types of social networks. Some networks, such as Facebook, model close relationships between people: a social friendship in which edges presumably form only between people who know each other personally. This is in direct contrast to systems such as Twitter and Google+, in which the act of “following” does not necessarily indicate such a connection. Presumably, a person trusts his or her friends but is interested in “following” not just friends but many others, including celebrities, political figures, acquaintances, corporations, and even enemies.

When agents participate in both networks, it is useful to be able to rank the nodes of a larger network based on the smaller. A node v can compute a ranking on the smaller, more closely knit, network and then use this to calculate a ranking of nodes in the larger network that do not appear in the more personal network. Current ranking mechanisms such as PageRank compute a global ranking and therefore are not suitable for this situation.

In this paper, we will show how Dirichlet PageRank can be used to model ranking problems, such as the above four, involving trust and distrust. We will give an efficient algorithm to compute Dirichlet PageRank approximately, which leads to efficient algorithms to solve these problems.

1.1. Related Work

The idea of ranking nodes in a graph has a rich history, beginning with the introduction of PageRank in [Brin and Page 98]. The original PageRank definition was designed for Web search, but many researchers have developed more tailored ranking systems such as personalized PageRank [Haveliwala 04, Jeh and Widom 03], which gives a ranking relative to some specified starting distribution \vec{s} .

One pitfall with PageRank as a ranking system is the fact that all edges contribute positively. In practice, an edge such as a link from one Web page to another can also represent a negative interaction or distrust between the nodes. Several related mathematical models of propagating trust and distrust in a network ranking system are given in [Guha et al. 04], and there are numerous empirical results. Another algorithm [Gyöngyi et al. 04] relies on a small hand-picked set of trusted nodes, but one must be careful not to allow malicious nodes to be included.

There are many other algorithms derived from PageRank that use specific heuristics to model trust or distrust in ranking schemes. For example, [Andersen et al. 08] considers axioms that a ranking system should satisfy and develops several ranking systems accordingly; [Borgs et al. 10] and [de Kerchove and Dooren 08] systematically model distrust by modifying the PageRank equations to consider negatively weighted edges; and [Kamvar et al. 03] gives an algorithm with a similar flavor using random walks. Many of these algorithms are closely related, but rigorous analysis is desired for capturing specific phenomena. We will show that these related models can be represented by Dirichlet PageRank with appropriate boundary conditions.

Another area of research concerns spam nodes when they are identified. It has been shown that if agents can collude [Baeza-Yates et al. 05] or easily create

pseudonyms [Cheng and Friedman 05], they can artificially boost their ranking in PageRank and other ranking systems. There has been some work done on how to effectively penalize these vertices [Benczur et al. 05], and our Dirichlet PageRank can be efficiently used to achieve the same goal.

1.2. Results in This Paper

Motivated by the continual development of new PageRank-based algorithms and the analysis of Dirichlet eigenvectors in [Chung and Yau 00], we develop and analyze Dirichlet PageRank vectors as well as an efficient algorithm to compute them. For a connected graph G , we give a Dirichlet PageRank equation and show how to compute the unique solution with Dirichlet boundary conditions: $\text{pr}(v) = 0$ for vertices v on the boundary of a specified vertex subset S .

After giving the algorithm for computing Dirichlet PageRank vectors, we generalize the boundary conditions to arbitrary values $\text{pr}(v) = \sigma(v)$ for boundary vertices v . We give an efficient algorithm, `ApproxDirichPR`, to compute approximate Dirichlet PageRank vectors with any boundary condition σ . We also give a full analysis leading to the following theorem. We use the notation $|\cdot|$ to denote the L_1 -norm, and for a subset S of vertices in G , the volume of S is denoted by $\text{vol}(S) = \sum_{v \in S} d_v$, where d_v is the degree of v in the graph G . Detailed definitions will be given in Section 2.

Theorem 1.5. *For any $\epsilon \in (0, 1)$ and any teleportation constant $\alpha \in (0, 1)$, the algorithm `ApproxDirichPR` outputs an ϵ -approximate Dirichlet PageRank vector $\tilde{\text{pr}}_S$ in time $O(\frac{\text{vol}(S) \log 1/\epsilon}{\alpha})$, which, compared to the exact Dirichlet PageRank pr_S , satisfies*

$$|\text{pr}_S - \tilde{\text{pr}}_S| < \frac{\epsilon \text{vol}(S)}{\alpha}.$$

We illustrate several applications of Dirichlet PageRank with boundary conditions below. Many of the specific PageRank variations are covered by this general framework, and we will show how its use can allow the efficient consideration of several models in [Andersen et al. 08, Benczur et al. 05, Borgs et al. 10].

1.3. Several Applications of Dirichlet PageRank

1.3.1. Allowing negative edges in the graph. While trust between two vertices is denoted by a positive weight, it is natural to quantify distrust as negative weights for the associated edges. There are many other types of relations in a network that can be represented with negative edges as well, and the usual PageRank vectors do

not consider negative weights. We will use Dirichlet PageRank as a tool to deal with graphs containing negative edges in Section 6.1.

1.3.2. Diminishing known spammers' influence. Many Web pages can be identified as spammers based on content or user reports. It is desirable to have a network-ranking scheme that takes such considerations into account by penalizing both the known spammer nodes and others with many links to them. We will show that Dirichlet PageRank is useful for dealing with spammer nodes in Section 6.2.

1.3.3. Considering trusted friends' opinions. A single node in a graph may have a set of trusted friends or neighbors whose opinions need to be considered strongly in designing a vertex-ranking scheme. If these trusted nodes have their own independent ranking opinions, we can use Dirichlet PageRank with appropriate boundary conditions to compute a trust-based ranking. We will give an algorithm, PRTrustedFriends, for this problem in Section 6.3.

1.3.4. Validating ranking for newly created nodes. Suppose that a new person enters a social network but is unsure about which nodes are trustworthy. Personalized PageRank is a useful tool for deriving quantitative information, but it raises the question whether this ranking is susceptible to unknown spammers. Without a specific set of trusted friends, it may seem hopeless for the newcomer, but Dirichlet PageRank with boundary conditions can be used to validate and adjust its ranking with a randomly selected pool of established nodes within the network. We will give the details for an algorithm, PRValidation, in Section 6.3.

1.3.5. Reconciling ranking in personal and global social networks. Several interesting questions arise in the analysis of different types of social networks. Some, like Facebook, offer a more personal viewpoint, as reflected in the network structure, whereby edges are formed only by mutual consent between two people who usually know each other. This is in contrast to a network such as Twitter, whose connections are often (though not always) impersonal. For example, people “follow” each other based not only on friendship, but also on subject matter, celebrity appeal, advertising, and many other reasons.

With the vast array of information available on a network such as Twitter, it is important for a user to know who is trustworthy or worth following. This is a difficult problem, but a user does have some information at hand, such as its own, more closely knit, smaller social networks or even a trusted subgraph of the larger network. Using Dirichlet PageRank, a user can compute a ranking on the smaller network, and then use boundary conditions appropriately to infer a ranking on the remaining nodes in the larger network, taking its personal

associations into account. We will develop an algorithm, PRTrustNetwork, for this problem in Section 6.3.

Finally, we can use similar ideas to tackle the problem in the reverse direction: Suppose a global ranking of the nodes in a large, loose social network such as Twitter is known, and a user wants to develop a personalized ranking for a small subgraph or its own trusted network, taking the global ranking into account. We again can use Dirichlet PageRank with appropriate boundary conditions, as outlined in the algorithm PRInferRanking, also in Section 6.3.

The rest of the paper proceeds as follows. In Section 2, we outline necessary background on PageRank and Dirichlet boundary conditions. Section 3 develops the theory of PageRank with Dirichlet boundary conditions, and Section 4 extends this theory for arbitrary boundary conditions σ . We develop and analyze the ApproxDirichPR algorithm in Section 5 and give algorithms for the previously discussed applications in Section 6.

2. Preliminaries

For a connected undirected graph $G = (V, E)$ with n vertices and m edges, let A be the *adjacency matrix* and D the *diagonal degree matrix*, where D_{ii} is the degree of the i th vertex. A typical random walk on G is defined by the *transition probability matrix* $D^{-1}A$. In this paper, we consider a *lazy random walk* that is defined by the *lazy transition probability matrix*

$$W = \frac{1}{2}(I + D^{-1}A).$$

The *normalized Laplacian* \mathcal{L} is defined by

$$\mathcal{L} = D^{-1/2}(D - A)D^{-1/2} = I - D^{-1/2}AD^{-1/2}.$$

The normalized Laplacian \mathcal{L} and its spectrum are useful in analyzing graphs and their associated random walks (see [Chung 97] for more details).

The *restricted Laplacian* \mathcal{L}_S is the submatrix of \mathcal{L} restricted to $S \times S$. The *restricted Green's function* $\mathcal{G}_{S,\beta}$ is defined by

$$\mathcal{G}_{S,\beta}(\beta I_S + \mathcal{L}_S) = I_S,$$

where $\beta \geq 0$. Note that \mathcal{L}_S is positive definite [Chung and Yau 00], so $\mathcal{G}_{S,\beta}$ is well defined.

The *PageRank vector* pr has two parameters: the *teleportation constant* $\alpha \geq 0$ and the *seed vector* s . For given α and s , the PageRank pr is defined to be the

unique solution to the PageRank equation

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

where we treat s and pr as row vectors. This paper uses the lazy random walk transition matrix W instead of the regular random walk transition matrix $D^{-1}A$, but the two PageRank definitions are equivalent up to a change in teleportation constant α (see [Andersen et al. 06]). PageRank was first introduced in [Brin and Page 98] to measure the importance of Web pages (with the seed vector $s = \vec{1}/n$), and it has since been applied to many problems, including the measurement of trust in social networks [Andersen et al. 08]. Fast approximation algorithms for PageRank can be found in [Andersen et al. 06, Chung and Yau 10].

We can also write PageRank as a geometric sum of random walks [Andersen et al. 06]:

$$\text{pr} = \alpha s + \alpha \sum_{t=1}^{\infty} (1 - \alpha)^t s W^t.$$

This allows us to see that pr , with the same teleportation constant α , is linear in its seed vector s .

3. PageRank with Dirichlet Boundary Conditions

Let S denote a subset of the vertex set V of G . The *volume* $\text{vol}(S)$ denotes the sum of the degrees of vertices in S . The *vertex boundary* $\delta(S)$ is defined by

$$\delta(S) = \{u \mid u \notin S \text{ and } \exists v \in S, (v, u) \in E\}.$$

The essential question proposed in Problem 1.1 is how to take into account the boundary edges. We remark that in the classical areas of differential equations defined on some geometric spaces, the boundary conditions refer to the constraints defined on the boundaries of specified regions and are imposed on the solutions of the equations. For the PageRank equation, the lazy random walk transition matrix W is closely related to the normalized Laplacian \mathcal{L} , which is the analogue of the Laplace–Bertrami operator in differential geometry. Thus, the PageRank equation (2.1) can be regarded as a discrete analogue of a set of differential equations, and Dirichlet PageRank can be viewed as a solution of the PageRank equation with boundary conditions. The basic problem of deriving PageRank vectors with Dirichlet boundary conditions was also previously examined in [Chung 10].

Let S be a subset of G . For a function (or vector) $f : V \rightarrow \mathbb{R}$, we say that f satisfies the *Dirichlet boundary condition* if

$$f(v) = 0 \text{ for all } v \in \delta(S).$$

The PageRank vector satisfying the Dirichlet boundary condition is the solution of the following equation, for all vertices v :

$$\text{pr}(v) = \begin{cases} \alpha s(v) + (1 - \alpha) \sum_{u \in V} \text{pr}(u) W_{uv} & \text{if } v \in S, \\ 0 & \text{otherwise.} \end{cases} \quad (3.1)$$

Let pr_S and s_S denote the vectors pr and s restricted to S , respectively. Similarly, let W_S, D_S, A_S denote the respective matrices restricted to $S \times S$.

Theorem 3.1. *For a connected graph G , vector s , and $\alpha > 0$, the PageRank equation (3.1) has one and only one solution. With $\beta = 2\alpha/(1 - \alpha)$, it is given by*

$$\text{pr}_S = \beta s_S D_S^{-1/2} \mathcal{G}_{S,\beta} D_S^{1/2}.$$

Proof. Since $\text{pr}(v) = 0$ when $v \notin S$, the PageRank equation (3.1) is equivalent to

$$\text{pr}_S = \alpha s_S + (1 - \alpha) \text{pr}_S W_S.$$

Since

$$W = \frac{1}{2}(I + D^{-1}A) = I - \frac{1}{2} \left(D^{-1/2} \mathcal{L} D^{1/2} \right)$$

and D is a diagonal matrix, we have

$$W_S = I_S - \frac{1}{2} \left(D_S^{-1/2} \mathcal{L}_S D_S^{1/2} \right).$$

Thus, we have

$$\text{pr}_S = \alpha s_S + (1 - \alpha) \text{pr}_S \left(I_S - \frac{1}{2} \left(D_S^{-1/2} \mathcal{L}_S D_S^{1/2} \right) \right).$$

Solving for pr_S yields the theorem; uniqueness follows from the uniqueness of $\mathcal{G}_{S,\beta}$. \square

Using Dirichlet PageRank vectors with Dirichlet boundary conditions, we can solve Problem 1.1 in a straightforward way, and it is clear that vertices outside of S will not influence the ranking. However, it is not immediately apparent how the boundary edges affect the ranking result. In Figure 1, comparisons are given between rankings obtained by three methods. For a small graph, we compute the Dirichlet PageRank vector and compare it with the following two alternative methods:




Figure 1. (a) Ranking computed by Method 1. (b) Ranking computed by Method 2. (c) Ranking computed by Dirichlet PageRank.

Method 1. Compute the PageRank for the entire graph and simply use these values on the subgraph.

Method 2. Delete the rest of the graph including boundary edges; then compute the PageRank for the remaining induced subgraph.

We note that the difference between Method 1 and Dirichlet PageRank is the relative ranking of vertices b and d . Using Method 1, b and d have the same ranking, while with Dirichlet PageRank, b has a higher ranking. This is consistent with the fact that in this subgraph, b is trusted by all other vertices, while d is trusted by only two out of three other nodes. Note that d is also trusted by another vertex outside of the subgraph, but the influence of the outside vertex is not as significant.

The difference between Method 2 and Dirichlet PageRank is the relative ranking of vertices c and d . By Method 2, c and d have the same ranking; by Dirichlet PageRank, d has a slightly higher ranking. This is consistent with the fact that c and d trust each other and are trusted by b . In addition, d is trusted by one vertex from outside, reflected by a higher ranking of d .

We give another Dirichlet PageRank example in Figure 2. In Figure 2(a), a social network [Lusseau et al. 03] has nodes shaded according to their PageRank (for the case of $\alpha = 0.1$). Now suppose we have identified two spammers and want to penalize their ranking and influence, as described in Problem 1.3. In Figure 2(b), we simply compute the usual PageRank and set the rank of the two spammers to zero. This is equivalent to the ranking algorithm proposed in [Gyöngyi et al. 04]. In Figure 2(c), we compute the Dirichlet PageRank with the boundary condition $\sigma(u) = \sigma(v) = 0$ for the spammers u and v . It is apparent




Figure 2. Three rankings of a network [Lusseau et al. 03]. Darker shades indicate higher ranking. (a) PageRank of the graph with $\alpha = 0.1$. (b) PageRank of the graph, with two “spammers” identified with zero rank. (c) Dirichlet PageRank of the graph, with the nodes around the two “spammers” showing decreased rank (color figure available online).

that the rankings of the nodes surrounding the spammers have been decreased, illustrating the effects of propagation of distrust.

These two examples serve as an illustration of the contributions of the boundary edges to the ranking, made more rigorous in the following lemma.

Let pr' be the PageRank vector computed by Method 1, and let pr'' denote the PageRank computed by Method 2. It is easy to see that for every $v \in S$, we have $\text{pr}(v) \leq \text{pr}'(v)$. We define two vertex sets S_o and S_i :

$$S_o = \{v \in S \mid \exists u \notin S : (u, v) \in E\} \quad \text{and} \quad S_i = S \setminus S_o.$$

Let W_{ii} and W_{oi} denote W restricted to $S_i \times S_i$ and $S_o \times S_i$, respectively. Similarly, we define two vectors w_o and w_i :

$$w_o = (1 - \alpha)\mathbf{1}W_{oi}^T \quad \text{and} \quad w_i = \mathbf{1} - (1 - \alpha)\mathbf{1}W_{ii}^T.$$

Lemma 3.2. *Suppose that both S_o and S_i are nonempty. Then we have*

$$\frac{\text{pr}'_{S_o} w_o^T}{\text{pr}'_{S_i} w_i^T} \geq \frac{\text{pr}_{S_o} w_o^T}{\text{pr}_{S_i} w_i^T} \geq \frac{\text{pr}''_{S_o} w_o^T}{\text{pr}''_{S_i} w_i^T}.$$

Proof. Let W'' denote the lazy random walk transition probability matrix of G_S , where G_S is the subgraph of G restricted to S . Then the following equations

hold:

$$\text{pr}'_{S_i} = \alpha s_{S_i} + (1 - \alpha) (\text{pr}'_{S_i} W_{ii} + \text{pr}'_{S_o} W_{oi}), \quad (3.2)$$

$$\text{pr}_{S_i} = \alpha s_{S_i} + (1 - \alpha) (\text{pr}_{S_i} W_{ii} + \text{pr}_{S_o} W_{oi}), \quad (3.3)$$

$$\text{pr}''_{S_i} = \alpha s_{S_i} + (1 - \alpha) (\text{pr}''_{S_i} W''_{ii} + \text{pr}''_{S_o} W''_{oi}). \quad (3.4)$$

Let $c_1 = \alpha s_{S_i} \mathbf{1}^T$; the definitions of w_o and w_i give

$$\frac{\text{pr}_{S_o} w_o^T}{\text{pr}_{S_i} w_i^T} = \frac{\text{pr}_{S_i} w_i^T - c_1}{\text{pr}_{S_i} w_i^T}.$$

Subtracting (3.3) from (3.2) yields

$$(\text{pr}'_{S_i} - \text{pr}_{S_i})(I - (1 - \alpha)W_{ii}) = (\text{pr}'_{S_o} - \text{pr}_{S_o})((1 - \alpha)W_{oi}).$$

Let $c_2 = (\text{pr}'_{S_i} - \text{pr}_{S_i})(I - (1 - \alpha)W_{ii}) \mathbf{1}^T$. Since for all $v \in S$, we have $\text{pr}(v) \leq \text{pr}'(v)$, it follows that $\text{pr}'_{S_i} - \text{pr}_{S_i}$ and c_2 are both nonnegative. Thus, we have

$$\frac{\text{pr}'_{S_o} w_o^T}{\text{pr}'_{S_i} w_i^T} = \frac{\text{pr}'_{S_i} w_i^T - c_1}{\text{pr}'_{S_i} w_i^T} = \frac{\text{pr}_{S_i} w_i^T + c_2 - c_1}{\text{pr}_{S_i} w_i^T + c_2}.$$

Since for every $v \in S_o$, the degree of v in G_S is strictly smaller than the degree of v in G , we have $W''_{v,u} \geq W_{v,u}(1 + 1/d)$, where d is the maximum degree among vertices in S_o in G_S . For $v \in S_i$, there is no change in degree from G to G_S , so $W_{ii} = W''_{ii}$. Hence, we have

$$\begin{aligned} \text{pr}''_{S_i} (I - (1 - \alpha)W_{ii}) &= \alpha s_{S_i} + \text{pr}''_{S_o} ((1 - \alpha)W''_{oi}) \\ &\geq \alpha s_{S_i} + \text{pr}''_{S_o} ((1 - \alpha)W_{oi}) \left(1 + \frac{1}{d}\right). \end{aligned}$$

Therefore,

$$\frac{\text{pr}''_{S_o} w_o^T}{\text{pr}''_{S_i} w_i^T} \leq \frac{\text{pr}_{S_i} w_i^T - c_1}{(1 + 1/d) \text{pr}_{S_i} w_i^T},$$

and the lemma follows from

$$\frac{\text{pr}_{S_i} w_i^T + c_2 - c_1}{\text{pr}_{S_i} w_i^T + c_2} \geq \frac{\text{pr}_{S_i} w_i^T - c_1}{\text{pr}_{S_i} w_i^T} \geq \frac{\text{pr}_{S_i} w_i^T - c_1}{(1 + 1/d) \text{pr}_{S_i} w_i^T}.$$

□

We remark that pr'' tends to underestimate the ranking, since it ignores all the boundary edges. In the other direction, the influence of boundary nodes should be taken into consideration but not so much as to overestimate the PageRank on S_o in comparison with pr'_{S_o} . Lemma 3.2 shows that pr_{S_o} is bounded between pr''_{S_o} and pr'_{S_o} , as desired.

4. Dirichlet PageRank with Given Boundary Conditions

In this section, we generalize Dirichlet PageRank to use arbitrary boundary conditions given by $\sigma : \delta(S) \rightarrow \mathbb{R}$. Note that σ can have negative values.

The Dirichlet PageRank vector with given boundary conditions σ is defined by the equations

$$\text{pr}(v) = \begin{cases} \alpha s(v) + (1 - \alpha) \sum_{u \in V} \text{pr}(u) W_{uv} & \text{if } v \in S, \\ \sigma(v) & \text{if } v \in \delta(S). \end{cases} \quad (4.1)$$

Here, we use the convention that $|\sigma| \leq 1$. Let $W_{\delta(S)}$ denote W restricted to $\delta(S) \times S$.

Theorem 4.1. *For a connected graph G , vector s , $\alpha > 0$, and given boundary conditions σ , the PageRank equation (4.1) has one and only one solution. With $\beta = 2\alpha/(1 - \alpha)$, it is given by*

$$\text{pr}_S = (\beta s_S + 2\sigma_{\delta(S)} W_{\delta(S)}) D_S^{-1/2} \mathcal{G}_{S,\beta} D_S^{1/2}.$$

Proof. Note that

$$\begin{aligned} \text{pr}_S &= \alpha s_S + (1 - \alpha) (\text{pr}_S W_S + \sigma_{\delta(S)} W_{\delta(S)}) \\ &= \frac{1 - \alpha}{2} (\beta s_S + 2\sigma_{\delta(S)} W_{\delta(S)}) + (1 - \alpha) \text{pr}_S W_S. \end{aligned} \quad (4.2)$$

The theorem follows by expressing W_S in terms of \mathcal{L}_S and solving for pr_S , as in the proof of Theorem 4.1. \square

To solve Problem 1.2, one can adjust the ranking by setting boundary conditions σ and solving the Dirichlet PageRank equation. We can use σ to specify known quantitative distrust, which will then propagate to the rest of vertices in S .

5. Algorithms and Analysis

Solving the PageRank equation (3.1) or (4.1) with boundary conditions requires both matrix–vector multiplication and solving a linear system of the form

$$x(\beta I_S + \mathcal{L}_S) = y.$$

The running time of solving the PageRank equation is dominated by the complexity of solving the linear system. Since the matrix $\beta I_S + \mathcal{L}_S$ is diagonally dominant, it can be solved approximately in nearly linear time with a Spielman–Teng solver [Spielman and Teng 08], but we will also give a simpler algorithm,

Algorithm 1. (ApproxDirichPR)

Input: $G, S, \alpha, s, \sigma, \epsilon$ **Output:** pr_S $\text{pr}_S \leftarrow \mathbf{0}, \epsilon' \leftarrow 1, r \leftarrow \alpha s_S + (1 - \alpha)\sigma_{\delta(S)}W_{\delta(S)}$ **while** $\epsilon' > \epsilon$ **do** **while** $|r(v)| \geq \epsilon' d_v$ for some v **do** $\text{pr}_S(v) \leftarrow \text{pr}_S(v) + r(v)$ For each neighbor u of v , $r(u) \leftarrow r(u) + (1 - \alpha)r(v)/2d_v$ $r(v) \leftarrow (1 - \alpha)r(v)/2$ **end while** $\epsilon' \leftarrow \epsilon'/2$ **end while**

ApproxDirichPR, to compute approximate Dirichlet PageRank vectors. This approximation algorithm is faster and has a better approximation ratio if the constant α is not too small.

The algorithm ApproxDirichPR (Algorithm 1) is outlined as follows: We initialize pr_S as $\mathbf{0}$ and maintain a residue r , which is the difference between the right side and left side of (4.2). Then we gradually move the ‘‘mass’’ from r to pr_S while maintaining the invariant

$$\text{pr}_S + r = \alpha s_S + (1 - \alpha) (\text{pr}_S W_S + \sigma_{\delta(S)} W_{\delta(S)})$$

until we have $r(v) \leq \epsilon' d_v$ for every $v \in S$. At the start, we set $\epsilon' = 1$. After each iteration, we decrease ϵ' by half until $\epsilon' \leq \epsilon$, which is the given desired approximation ratio.

To analyze the above algorithm, the proof of Theorem 1.5 is given as follows.

Proof of Theorem 1.5. To bound the running time, we first show that in each iteration of the inner loop, $|r|_1$ will decrease by at least $\alpha \epsilon' d_v$. Let r^b be r before the iteration and let r^a be r after the iteration. We have

$$\begin{aligned} |r^a|_1 &= |r^a(v)| + \sum_{u \neq v} |r^a(u)| = \frac{1 - \alpha}{2} |r^b(v)| + \sum_{u \neq v} \left| r^b(u) + \frac{1 - \alpha}{2d_v} r^b(v) \right| \\ &\leq \frac{1 - \alpha}{2} |r^b(v)| + \sum_{u \neq v} \left(|r^b(u)| + \frac{1 - \alpha}{2d_v} |r^b(v)| \right) \\ &= (1 - \alpha) |r^b(v)| + \sum_{u \neq v} |r^b(u)| = |r^b|_1 - \alpha |r^b(v)|. \end{aligned}$$

We note that the above equations hold for both positive and negative values of $r^b(v)$. Since v is chosen to satisfy $|r(v)| \geq \epsilon' d_v$, it follows that $|r^a|_1 \leq |r^b|_1 - \alpha \epsilon' d_v$.

Note that at the beginning of each iteration of the outer loop, we have $|r| \leq 2\epsilon' \text{vol}(S)$. Let T be the number of iterations of the inner loop and v_i the vertex selected at the i th iteration for $1 \leq i \leq T$. We have

$$\sum_{i=1}^T \alpha \epsilon' d_{v_i} \leq 2\epsilon' \text{vol}(S),$$

which implies

$$\sum_{i=1}^T d_{v_i} \leq \frac{2 \text{vol}(S)}{\alpha}.$$

Since a FIFO queue can be used to store every vertex v such that $|r(v)| \geq \epsilon' d_v$, each iteration of the inner loop can be completed in $O(d_v)$ time. Therefore, the running time of one outer iteration is $\sum_{i=1}^T d_{v_i}$, which is bounded from above by $2 \text{vol}(S)/\alpha$.

There are $\log 1/\epsilon$ iterations of the outer loop; therefore, the overall running time is

$$O\left(\frac{\text{vol}(S) \log 1/\epsilon}{\alpha}\right).$$

To prove the correctness of the approximation ratio, we will first show that the following invariant is maintained during the entire algorithm:

$$\text{pr}_S + r = \alpha s_S + (1 - \alpha) (\text{pr}_S W_S + \sigma_{\delta(S)} W_{\delta(S)}).$$

This equation holds trivially in the beginning, where $\text{pr}_S = \mathbf{0}$ and $r = \alpha s_S + (1 - \alpha) \sigma_{\delta(S)} W_{\delta(S)}$. For each inner iteration, let r^b , pr_S^b be r , pr_S before the iteration, and let r^a , pr_S^a be r , pr_S after the iteration. We have

$$\begin{aligned} & \text{pr}_S^a(v) + r^a(v) \\ &= \text{pr}_S^b(v) + r^b(v) + \frac{1 - \alpha}{2} r^b(v) \\ &= \alpha s_S(v) + (1 - \alpha) ([\text{pr}_S^b W_S](v) + [\sigma_{\delta(S)} W_{\delta(S)}](v)) + \frac{1 - \alpha}{2} r^b(v) \\ &= \alpha s_S(v) \\ & \quad + (1 - \alpha) \left(\frac{1}{2} \text{pr}_S^b(v) + \sum_{w \neq v} \frac{1}{2d_w} \text{pr}_S^b(w) + [\sigma_{\delta(S)} W_{\delta(S)}](v) \right) + \frac{1 - \alpha}{2} r^b(v) \end{aligned}$$

$$\begin{aligned}
&= \alpha s_S(v) \\
&\quad + (1 - \alpha) \left(\frac{1}{2} (\text{pr}_S^b(v) + r^b(v)) + \sum_{w \neq v} \frac{1}{2d_w} \text{pr}_S^b(w) + [\sigma_{\delta(S)} W_{\delta(S)}](v) \right) \\
&= \alpha s_S(v) + (1 - \alpha) \left(\frac{1}{2} \text{pr}_S^a(v) + \sum_{w \neq v} \frac{1}{2d_w} \text{pr}_S^a(w) + [\sigma_{\delta(S)} W_{\delta(S)}](v) \right) \\
&= \alpha s_S(v) + (1 - \alpha) ([\text{pr}_S^a W_S](v) + [\sigma_{\delta(S)} W_{\delta(S)}](v)),
\end{aligned}$$

and for $u \neq v$, we have

$$\begin{aligned}
&\text{pr}_S^a(u) + r^a(u) \\
&= \text{pr}_S^b(u) + r^b(u) + \frac{1 - \alpha}{2d_v} r^b(v) \\
&= \alpha s_S(u) + (1 - \alpha) ([\text{pr}_S^b W_S](u) + [\sigma_{\delta(S)} W_{\delta(S)}](u)) + \frac{1 - \alpha}{2d_v} r^b(v) \\
&= \alpha s_S(u) \\
&\quad + (1 - \alpha) \left(\frac{1}{2} \text{pr}_S^b(u) + \sum_{w \neq u} \frac{1}{2d_w} \text{pr}_S^b(w) + [\sigma_{\delta(S)} W_{\delta(S)}](u) \right) + \frac{1 - \alpha}{2d_v} r^b(v) \\
&= \alpha s_S(u) \\
&\quad + (1 - \alpha) \left(\frac{1}{2} \text{pr}_S^a(u) + \sum_{w \neq u} \frac{1}{2d_w} \text{pr}_S^a(w) - \frac{1}{2d_v} r^b(v) + [\sigma_{\delta(S)} W_{\delta(S)}](u) \right) \\
&\quad + \frac{1 - \alpha}{2d_v} r^b(v) \\
&= \alpha s_S(u) + (1 - \alpha) \left(\frac{1}{2} \text{pr}_S^a(u) + \sum_{w \neq u} \frac{1}{2d_w} \text{pr}_S^a(w) + [\sigma_{\delta(S)} W_{\delta(S)}](u) \right) \\
&= \alpha s_S(u) + (1 - \alpha) ([\text{pr}_S^a W_S](u) + [\sigma_{\delta(S)} W_{\delta(S)}](u)).
\end{aligned}$$

Thus, the invariant is maintained during the entire algorithm. Note that the equations hold for both positive and negative values of r . As a result, the output $\tilde{\text{pr}}_S$ satisfies

$$\tilde{\text{pr}}_S + r = \alpha s_S + (1 - \alpha) (\tilde{\text{pr}}_S W_S + \sigma_{\delta(S)} W_{\delta(S)}),$$

where $|r(v)| < \epsilon d_v$ for all vertices $v \in S$, and the exact solution pr_S satisfies

$$\text{pr}_S = \alpha s_S + (1 - \alpha) (\text{pr}_S W_S + \sigma_{\delta(S)} W_{\delta(S)}).$$

Taking the difference of these two equations, we get

$$\text{pr}_S - \tilde{\text{pr}}_S = r + (1 - \alpha)((\text{pr}_S - \tilde{\text{pr}}_S)W_S).$$

Since

$$|(\text{pr}_S - \tilde{\text{pr}}_S)W_S|_1 \leq |\text{pr}_S - \tilde{\text{pr}}_S|_1,$$

we have

$$|\text{pr}_S - \tilde{\text{pr}}_S|_1 \leq |r|_1 + (1 - \alpha)|(\text{pr}_S - \tilde{\text{pr}}_S)W_S|_1 \leq |r|_1 + (1 - \alpha)|\text{pr}_S - \tilde{\text{pr}}_S|_1,$$

which implies

$$|\text{pr}_S - \tilde{\text{pr}}_S|_1 \leq \frac{1}{\alpha}|r|_1 < \frac{\epsilon \text{vol}(S)}{\alpha}.$$

□

6. Applications of Dirichlet PageRank

6.1. Allowing Negative Edges in the Graph

Negative edges arise in many network problems, making PageRank less suitable for finding a desirable ranking (see [de Kerchove and Dooren 08]). There have been numerous attempts to address this problem. One way is to ignore the entries corresponding to the negative links, as seen as in [Kamvar et al. 03, Langville and Meyer 06]. By treating a negative link between two vertices the same as no link [Kamvar et al. 03, Langville and Meyer 06], the PageRank vector can be computed and used as a ranking. Unfortunately, the information contained in those negative links is lost. To remedy this, in [Guha et al. 04], a trust ranking is computed based on only positive links and one single step of propagation of distrust. Namely, the distrust is propagated only to immediate neighbors without influencing the rest of the vertices. A more sophisticated algorithm, PageTrust, is proposed in [de Kerchove and Dooren 08], which uses a fairly complicated update rule relying on a relatively large number of iterations until convergence. However, the running time for each iteration is quite large: $O(\bar{d}nn^-)$, where \bar{d} is the average degree and n^- is the number of vertices receiving negative links. Thus, the worst-case complexity is $O(n^3)$.

Using Dirichlet PageRank, we develop a simple and fast algorithm to propagate distrust in graphs with negative edges. The key idea can be outlined as follows: we first compute the usual PageRank based on positive edges; then, based on the ranking result, we convert negative links to boundary conditions and then compute the Dirichlet PageRank.

Algorithm 2. (NegLinkPageRank)

Input: $G = (V, E)$, v, α, ϵ **Output:** prDetermine \widehat{G} , V^s and E^+ as described above.pr⁺ \leftarrow SharpApproximatePR(v, α, ϵ) using the graph (V, E^+) $\sigma(u) \leftarrow \frac{d^-}{d^+}$ pr⁺(u) for each vertex u in V^s pr \leftarrow ApproxDirichPR($\widehat{G}, V, \alpha, v, \sigma, \epsilon$)

Let E^+ be the set of positive edges, E^- the set of negative edges, and V^- the set of vertices incident to negative edges. Let $d^-(u)$ and $d^+(u)$ denote the numbers of negative and positive edges incident to u , respectively. For each vertex $u \in V^-$, we create a *shadow vertex* u^s . Let V^s denote the set of all shadow vertices, and define

$$E^s = \{\{u^s, v\}, \{u, v^s\} \mid \{u, v\} \in E^-\}.$$

We then form a new graph $\widehat{G} = (\widehat{V}, \widehat{E})$, where $\widehat{V} = V \cup V^s$ and $\widehat{E} = E^+ \cup E^s$. In Algorithm 2, we will set boundary conditions on V^s in \widehat{G} to propagate distrust.

The intuition behind setting the values of the boundary condition $\sigma(u)$ is as follows: We let $\sigma(u)/d^- = \text{pr}^+(u)/d^+$. Namely, for a vertex, the amount of distrust propagated via the negative edge is equal to the amount of trust propagated via the positive edge. The running time of this algorithm is nearly linear using our ApproxDirichPR algorithm.

As an example of using Dirichlet PageRank on a graph with positive and negative edges, we examine a network of tribes in New Guinea studied in the mid-twentieth century [Hage and Harary 83, Read 54]. A positive edge indicates a tribal alliance, and negative edges represent enemy relationships. As illustrated in Figure 3(a), the positive edges are in light green, and the negative edges are in light red. One way to calculate a vertex ranking is to ignore the negative edges, as shown in Figure 3(a). The PageRank vector computed in this manner is actually the uniform distribution. In Figure 3(b), we use Dirichlet PageRank to compute a vertex ranking taking the negative edges into account. It is apparent that vertices are appropriately ranked by taking advantages of their trusting and distrusting relationships.

6.2. Adjusting Spammers' Influence

One disadvantage of purely link-based ranking systems such as PageRank is that they interpret all nodes as honest agents and all links as votes or validation




Figure 3. Comparison of vertex rankings in a network [Hage and Harary 83, Read 54]. (a) PageRank, computed by ignoring negative links. (b) Dirichlet PageRank, taking negative links into account (color figure available online).

between nodes. However, real-world networks such as the World Wide Web often contain malicious nodes or spammers. Of interest is to find ranking systems that can better represent the true ranking of nodes in the graph.

There are many schemes that have been developed to tackle such problems [Andersen et al. 08, Benczur et al. 05, Borgs et al. 10, Cheng and Friedman 05, Guha et al. 04, Gyöngyi et al. 04, Kamvar et al. 03, de Kerchove and Dooren 08]. It turns out that many of them can be modeled using Dirichlet PageRank with different boundary conditions. For example, [Benczur et al. 05] outlines an algorithm, SpamRank, that penalizes spam nodes. This algorithm uses sampling to find nodes whose PageRank vectors are significantly different from their neighbors', and gives a heuristic penalty score for each node. It then uses these penalty scores as a seed vector for personalized PageRank computation. However, if the penalties are represented by a probability distribution, the sampling techniques can be problematic for vertices with low degree. Using Dirichlet PageRank, we can penalize known spammers v by enforcing the condition $\text{pr}(v) = 0$. This is done in the example given in Figure 2. Furthermore, one can adjust the ranking even further by enforcing $\text{pr}(v) = -1$.

In [Borgs et al. 10], the trust and distrust are propagated within a network through a weighted random walk W with a trusted seed vertex s by assigning at the start, the rank $\text{pr}(s) = 1$. This can be regarded as Dirichlet PageRank with the boundary condition $\sigma(s) = 1$. There is a subtle difference in the way

Algorithm 3. (PRTrustedFriends)

Input: $G = (V, E)$, v , α , ρ , ϵ **Output:** \vec{p} Compute v 's ranking: $\vec{p} \leftarrow \text{SharpApproximatePR}(v, \alpha, \epsilon)$ [Chung and Yau 10]Compute v 's neighbors' rankings: for each neighbor u , $\vec{p}_u \leftarrow \text{SharpApproximatePR}(u, \alpha, \epsilon)$ Take a weighted average of v 's neighbors' rankings: $\vec{p}' \leftarrow \sum_{u \sim v} p(u) / \sum_{u \sim v} p(u) \vec{p}_u$.Take a set S of nodes that v ranks highly: $S \leftarrow \arg \max_{S \subseteq V, |S| \leq \rho |V|} \sum_{s \in S} p(s)$ Use v 's friends' ranking of S to adjust \vec{p} : $\vec{p} \leftarrow \text{ApproxDirichPR}(G, V \setminus S, \alpha, v, \vec{p}', \epsilon)$

distrust is handled (the algorithm in [Borgs et al. 10] does not allow for the propagation of trust scores less than 0). We note that Dirichlet PageRank allows us to efficiently consider these and many other models.

6.3. Adjusting Rank Based on Trust

While it is important to devise ranking systems that take known spammers into account, it is also crucial to be able to calculate a ranking based on various notions of trust in a network. There are numerous scenarios in which Dirichlet PageRank with boundary conditions can serve as a useful algorithmic tool.

We consider the following problem: In a network G , the node v wants to compute a personalized ranking of the nodes, but v trusts its own friends and wants its ranking in the top ρ fraction of nodes to be similar to its friends' rankings. This quantifies the assumptions that one's friends' actions carry a great deal of weight in one's own decisions. Vertex v can efficiently compute a personalized PageRank vector as its ranking function using algorithms from [Chung and Yau 10]. However, PageRank alone will not take into account the implied trust between v and its friends. But using Dirichlet PageRank with boundary conditions, we can take v 's trusted friends into account. We illustrate this in the algorithm PRTrustedFriends (Algorithm 3).

A natural extension of PRTrustedFriends can be described as the case that v is a newcomer to a network and is therefore unsure about what other nodes are trustworthy. In such a scenario, the only available information to v is the network itself. For ranking purposes, v can select a small number of nodes to compare with its own ranking. If these nodes are well distributed, this provides some control to ensure that v 's own ranking function is not too distorted by the presence of nearby spammers or malicious nodes. We give the algorithm PRValidation as Algorithm 4.

Algorithm 4. (PRValidation)

Input: $G = (V, E)$, v , k , α , ρ , ϵ **Output:** \vec{p} Compute v 's ranking: $\vec{p} \leftarrow \text{SharpApproximatePR}(v, \alpha, \epsilon)$ [Chung and Yau 10] $v_1, \dots, v_k \leftarrow$ i.i.d. samples from V according to \vec{p}

Compute rankings for the sampled nodes:

 $\vec{p}_k \leftarrow \text{SharpApproximatePR}(v_k, \alpha, \epsilon)$

Take a weighted average of these sampled rankings:

$$\vec{p}' \leftarrow \frac{1}{\sum_{i=1}^k p(v_k)} \sum_{i=1}^k p(v_k) \vec{p}_k$$

Take a set S of nodes that v ranked highly:

$$S \leftarrow \arg \max_{S \subseteq V, |S| \leq \rho |V|} \sum_{s \in S} p(s)$$

Use the sampled rankings to adjust \vec{p} :

$$\vec{p} \leftarrow \text{ApproxDirichPR}(G, V \setminus S, \alpha, v, \vec{p}', \epsilon)$$

A third, more complex, situation arises in the context of different types of social networks. Although the setup here appears somewhat complicated, it is a natural model for a common social phenomenon, addressing distinctions among different types of social networks.

Suppose that we have two networks $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ with $V_1 \subseteq V_2$. We interpret G_1 as a closely knit social network in which edges represent a deep mutual trust for one another. The network G_2 is a larger network whose edges are formed for relatively weak reasons, such as acquaintance or curiosity. We assume that a vertex v does not know much about the many sources of information acquired through G_2 . An important question for v is to determine which nodes in G_2 are trustworthy. How should the vertices of G_2 be ranked by taking advantage of G_1 ?

One effective way of finding such a ranking for vertices v in G_2 is to compute the ranking on G_1 first and then compute the Dirichlet PageRank on G_2 using G_1 's ranking as the boundary condition. This is outlined in the algorithm PRTrustNetwork (Algorithm 5).

Dirichlet PageRank can also be used to solve a related problem when a global ranking for G_2 is already known or precomputed. Suppose that such a ranking for G_2 exists, and a vertex $v \in G_1$ wishes to be able to rank G_1 by taking this into account. One way to do this is to compute a Dirichlet PageRank vector for

Algorithm 5. (PRTrustNetwork)

Input: $G_1 = (V_1, E_1)$, $G_2 = (V_2, E_2)$, $v \in V_1 \cap V_2$, α , ϵ **Output:** \vec{q} $\vec{p} \leftarrow \text{SharpApproximatePR}(v, \alpha, \epsilon)$ [Chung and Yau 10] for G_1 $\vec{q} \leftarrow \text{ApproxDirichPR}(G_2, V_2 \setminus V_1, \alpha, v, \vec{p}, \epsilon)$

Algorithm 6. (PRInferRanking)

Input: $G_2 = (V_2, E_2)$, $G_1 = (V_1, E_1) \subseteq G_2$, $v \in V_1$, \vec{p} , α , ϵ **Output:** \vec{q} $\partial E_1 \leftarrow \{(w, x) \in E_2 \mid w \in V_1, x \notin V_1\}$ $\partial V_1 \leftarrow \{w \in V_2 \setminus V_1 \mid w \text{ is an endpoint of an } e \in \partial E_1\}$ $\vec{q} \leftarrow \text{ApproxDirichPR}((V_1 \cup \partial V_1, E_1 \cup \partial E_1), V_1, \alpha, v, \vec{p}, \epsilon)$

G_1 but using the adjacent nodes in G_1 as a boundary with ranking given by the global ranking on G_2 . This procedure is given in the algorithm PRInferRanking (Algorithm 6).

In this section, we have examined several examples as applications of Dirichlet PageRank vectors. The list is by no means complete. These examples offer a glimpse of the applicability and flexibility of Dirichlet PageRank. Further applications and research directions remain to be explored.

Acknowledgments. This research is partially supported by ONR MURI N000140810747 and AFSOR 55082.

References

- [Andersen et al. 06] R. Andersen, F. Chung, and K. Lang. “Local Graph Partitioning Using PageRank Vectors.” In *FOCS*, 2006.
- [Andersen et al. 08] R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. “Trust-Based Recommendation Systems: An Axiomatic Approach.” In *WWW*, 2008.
- [Baeza-Yates et al. 05] R. Baeza-Yates, C. Castillo, and V. López. “PageRank Increase under Different Collusion Topologies.” In *Proceedings of the 1st International Workshop on Adversarial Information Retrieval on the Web*, 2005.
- [Benczur et al. 05] A. Benczur, K. Csalogany, T. Sarlos, and M. Uher. “SpamRank—Fully Automatic Link Spam Detection.” In *Proceedings of the 1st International Workshop on Adversarial Information Retrieval on the Web*, 2005.
- [Borgs et al. 10] C. Borgs, J. Chayes, A. T. Kalai, A. Malekian, and M. Tennenholtz. “A Novel Approach To Propagating Distrust.” In *WINE*, 2010.
- [Brin and Page 98] S. Brin and L. Page. “The Anatomy of a Large-Scale Hypertextual Web Search Engine.” *Computer Networks and ISDN Systems* 30:1–7 (1998), 107–117.
- [Cheng and Friedman 05] A. Cheng and E. Friedman. “Sybilproof Reputation Mechanisms.” In *Proceedings of Third Workshop on Economics of Peer-to-Peer Systems*, 2005.
- [Chung 97] F. Chung. *Spectral Graph Theory*. AMS Publications, 1997.
- [Chung 10] F. Chung. “PageRank as a Discrete Green’s Function.” *Geometry and Analysis I, ALM* 17 (2010), 285–302.

- [Chung and Yau 00] F. Chung and S.-T. Yau. “Discrete Green’s Functions.” *J. Combinatorial Theory (A)* 91 (2000), 191–214.
- [Chung and Yau 10] F. Chung and W. Zhao. “A Sharp PageRank Algorithm with Applications to Edge Ranking and Graph Sparsification.” *Proceedings of Workshop on Algorithms and Models for the Web Graph (WAW 2010)*, Lecture Notes in Computer Science 6516, pp. 2–14. Springer, 2010.
- [Guha et al. 04] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. “Propagation of Trust and Distrust.” In *WWW*, 2004.
- [Gyöngyi et al. 04] Z. Gyöngyi, H. Garcia-Molina, and J. Pedersen. “Combating Web Spam with TrustRank.” In *VLDB*, 2004.
- [Hage and Harary 83] P. Hage and F. Harary. *Structural Models in Anthropology*. Cambridge University Press, 1983.
- [Haveliwala 04] T. Haveliwala. “Topic-Sensitive Pagerank: A Context-Sensitive Ranking Algorithm for Web Search.” *IEEE Transactions on Knowledge and Data Engineering* 15 (2004), 784–796.
- [Jeh and Widom 03] G. Jeh and J. Widom. “Scaling Personalized Web Search.” In *WWW*, 2003.
- [Kamvar et al. 03] S. Kamvar, M. Schlosser, and H. Garcia-Molina. “The EigenTrust Algorithm for Reputation Management in P2P Networks.” In *WWW*, 2003.
- [de Kerchove and Dooren 08] C. de Kerchove and P. Dooren. “The PageTrust Algorithm: How to Rank Web Pages When Negative Links Are Allowed?” In *Proceedings of the SIAM International Conference on Data Mining*, 2008.
- [Langville and Meyer 06] A. Langville and C. Meyer. *Google’s PageRank and Beyond: The Science of Search Engine Rankings*. Princeton University Press, 2006.
- [Lusseau et al. 03] D. Lusseau, K. Schneider, O. J. Boisseau, P. Haase, E. Slooten, and S. M. Dawson. “The Bottlenose Dolphin Community of Doubtful Sound Features a Large Proportion of Long-Lasting Associations.” *Behavioral Ecology and Sociobiology* 54:4 (2003), 396–405.
- [Read 54] K. Read. “Cultures of the Central Highlands, New Guinea.” *Southwestern Journal of Anthropology* 10 (1954), 1–43.
- [Spielman and Teng 08] D. A. Spielman and S.-H. Teng. “Nearly-Linear Time Algorithms for Preconditioning and Solving Symmetric, Diagonally Dominant Linear Systems.” Available at <http://arxiv.org/abs/cs.NA/0607105>, 2008.

Fan Chung, Department of Computer Science and Engineering, University of California, San Diego, San Diego, CA 92093 (fan@cs.ucsd.edu)

Alexander Tsiasas, Department of Computer Science and Engineering, University of California, San Diego, San Diego, CA 92093 (atsiasas@cs.ucsd.edu)

Wensong Xu, Department of Computer Science and Engineering, University of California, San Diego, San Diego, CA 92093 (w4xu@cs.ucsd.edu)