

On Local Estimations of PageRank: A Mean Field Approach

Santo Fortunato, Marián Boguñá, Alessandro Flammini, and Filippo Menczer

Abstract. PageRank is a key element in the success of search engines, allowing the display of the most relevant hits in the first screen of results. One key aspect that distinguishes PageRank from other prestige measures such as in-degree is its global nature. From the information provider perspective, this makes it difficult or even impossible to predict how their pages will be ranked. Consequently, a market has emerged for the optimization of search engine results. Here we study the accuracy with which PageRank can be approximated by in-degree, a local measure made freely available by search engines. Theoretical and empirical analyses lead us to conclude that, given the weak degree of correlations in the Web link graph, the approximation can be relatively accurate, giving service and information providers an effective new marketing tool.

I. Introduction

PageRank has played a key role in the success of Web search engines, allowing them to display the most relevant hits in the first page of results. While recent results suggest that other query-independent criteria may be equally as effective as or more effective than PageRank [Richardson et al. 06], the introduction of PageRank as a factor in sorting results [Brin and Page 98] has contributed considerably to Google's lasting dominance in the search engine market [Sullivan 06].

PageRank is not the only possible measure of importance or prestige among Web pages. The simplest way to measure the prestige of a page is to count the incoming links (in-links) to the page. The number of in-links (in-degree)

is the number of citations that a page receives from other pages, so there is a correlation between in-degree and quality, especially when the in-degree is large. The in-degree of Web pages is very cheap to compute and maintain, so a search engine can easily keep in-degree updated with the evolution of the Web.

In-degree is a local measure, however. All links to a page are considered equal, regardless of where they come from. Two pages with the same in-degree are considered equally important, even if one is cited by much more prestigious sources than the other. To modulate the prestige of a page with that of the pages pointing to it means to move from the examination of an individual node in the link graph to that of the node together with its predecessor neighbors. PageRank represents such a shift from the local measure given by in-degree toward a global measure, where each Web page contributes to the importance of every other page.

The use of PageRank in place of in-degree for applications such as ranking by Web search engines relies on two assumptions:

1. PageRank is a quantitatively different and better prestige measure than in-degree.
2. PageRank cannot be easily guessed or approximated by in-degree.

There seems to be plenty of anecdotal and indirect evidence in support of these assumptions—for example, the popularity of PageRank—but little quantitative data to validate them. To wit, Amento et al. [Amento et al. 00] report a high average correlation between in-degree and PageRank (Spearman $\rho = 0.93$, Kendall $\tau = 0.83$). Further, they report the same average precision at 10 (60%). Similar findings about high correlation and comparable performance between PageRank and in-degree are also reported for specific tasks of retrieving home pages [Upstill et al. 03b] and corporations [Upstill et al. 03a]. These studies are based on relevance assessments by human subjects for relatively small numbers of queries. In this paper we take a more direct and quantitative approach in exploring the above PageRank assumptions by answering the following questions: What is the correlation between in-degree and PageRank across representative samples of the Web? How accurately can one approximate PageRank from local knowledge of in-degree?

From the definition of PageRank, other things being equal, the PageRank of a page grows with the in-degree of the page. Beyond this zero-order approximation, the actual relation between PageRank and in-degree has not yet been thoroughly investigated. It is known that the distributions of PageRank and in-degree follow an almost identical pattern [Pandurangan et al. 02, Donato et al. 04], i.e., a curve ending with a broad tail that follows a power law with exponent $\gamma \simeq 2.1$.

While Litvak et al. [Litvak et al. 07] provide a theoretical interpretation of this observation, the equal distributional behavior of PageRank and in-degree may indicate a strong correlation between the two variables. Surprisingly there is no agreement in prior literature about the correlation between PageRank and in-degree. Pandurangan et al. [Pandurangan et al. 02] show a weak correlation based on analysis of the brown.edu domain and the TREC WT10g collection. Donato et al. [Donato et al. 04] report on a correlation coefficient that is basically 0 based on analysis of a much larger sample ($2 \cdot 10^8$ pages) taken from the WebBase [Hirai et al. 00] collaboration. On the other hand, analysis of the University of Notre Dame domain by Nakamura [Nakamura 03] reveals a strong correlation.

In Section 2, we estimate PageRank for a generic directed network within a mean field approach [Binney et al. 92], which has been applied successfully to describe several other complex networks [Pastor-Satorras and Vespignani 04]. We obtain a system of self-consistent relations for the average value of PageRank of all vertices with equal in-degree. For a network without degree-degree correlations, the average PageRank turns out to be proportional to the in-degree, modulo an additive constant.

The prediction is validated empirically in Section 3, where we solve the equations numerically for four large samples of the Web graph; in each case the agreement between our theoretical estimate and the empirical data is excellent. We find that the Web graph is basically uncorrelated, so the average PageRank for each degree class can be well approximated by a linear function of the in-degree. As an additional contribution we settle the issue of the correlation between PageRank and in-degree; the linear correlation coefficient is consistently large for all four samples we have examined, in agreement with Nakamura [Nakamura 03]. We also calculate the size of the fluctuations of PageRank about the average value and find that the relative fluctuations decrease as the in-degree increases, indicating that our mean field estimate becomes more accurate for important pages.

Our results suggest that we can approximate PageRank from in-degree. By deriving PageRank with our formula we can predict the rank of a page within a hit list by knowing its in-degree and the number of hits in the list. Section 4 reports on an empirical study of this prediction, performed by submitting AltaVista queries to Google. The actual ranks turn out to be scattered about the corresponding predictions. The implication of the fact that PageRank is mostly determined by in-degree is that it is possible to estimate the number of in-links that a new page needs in order to achieve a certain rank among all pages dealing with a specific topic. This provides search marketers and information providers with a new, powerful tool to guide their campaigns.

2. Theoretical Analysis

The PageRank $p(i)$ of a page i is defined through the following expression [Brin and Page 98]:

$$p(i) = \frac{q}{N} + (1 - q) \sum_{j:j \rightarrow i} p(j)/k_{\text{out}}(j), \quad i = 1, 2, \dots, N, \quad (2.1)$$

where N is the total number of pages, $j \rightarrow i$ indicates a hyperlink from j to i , $k_{\text{out}}(j)$ is the out-degree of page j , and $1 - q$ is the so-called damping factor. The set of Equations (2.1) can be solved iteratively. From Equation (2.1) it is clear that the PageRank of a page grows with the PageRank of the pages that point to it. The sum over predecessor neighbors implies, however, that PageRank also increases with the in-degree of the page.

PageRank can be thought of as the stationary probability of a random walk process with additional random jumps. The steady state can be derived from the concentration of a large number of non-interacting random walkers. The physical description of the process is as follows: when a random walker is in a node of the network, at the next time step it jumps to a randomly chosen node with probability q and with probability $1 - q$ it moves to one of its successors with uniform probability. In the case of directed networks, there is the possibility that the node has no successors. In this case, the walker jumps to a randomly chosen node of the network with probability 1. The PageRank of a node i , $p(i)$, is the probability of finding the walker at node i when the process has reached the steady state, a condition that is always guaranteed by the jumping probability q .

The probability of finding the walker at node i at time step n follows a simple Markovian equation:

$$p_n(i) = \frac{q}{N} + (1 - q) \sum_{j:k_{\text{out}}(j) \neq 0} \frac{a_{ji}}{k_{\text{out}}(j)} p_{n-1}(j) + \frac{1 - q}{N} \sum_{j:k_{\text{out}}(j) = 0} p_{n-1}(j), \quad (2.2)$$

where a_{ji} is the adjacency matrix with entry 1 if there is a direct connection between j and i and 0 otherwise. The first term in Equation (2.2) is the contribution of walkers that decide to jump to a randomly chosen node, the second term is the random walk contribution, and the third term accounts for walkers that at the previous step were in dangling points and now jump to random nodes. In the limit $n \rightarrow \infty$, this last contribution becomes a constant term affecting all the nodes in the same way, and thus it can be removed from Equation (2.2) under the constraint that the final solution is properly normalized. Strictly speaking, this would lead to an effective teleportation term, which we omit to keep the notation simple. Alternatively, dangling nodes could be taken into account by a

proper rescaling of the second term [Del Corso et al. 05]. Hereafter, we intend all sums over nodes to exclude dangling ends, considering only nodes with $k_{\text{out}} > 0$.

The PageRank of page i is the steady state solution of Equation (2.2), $p(i) = \lim_{n \rightarrow \infty} p_n(i)$. Equation (2.2) can be used as a numerical algorithm to compute PageRank but, unfortunately, it is not possible to extract analytical solutions from it. In the next subsection we propose a mean field solution of Equation (2.2) that, nevertheless, gives a very accurate description of the PageRank structure of the Web. The mean field approach is often used in statistical physics and is reliable when each element of the system has many interaction partners.¹ In such cases the effect of the interactions can be taken into account in an averaged way, neglecting the variations among the elements.

2.1. Mean Field Analysis

Instead of analyzing the PageRank of single pages, we aggregate pages in classes according to their degree $\mathbf{k} \equiv (k_{\text{in}}, k_{\text{out}})$ and define the average PageRank of nodes of degree class \mathbf{k} as

$$\bar{p}_n(\mathbf{k}) \equiv \frac{1}{NP(\mathbf{k})} \sum_{i \in \mathbf{k}} p_n(i). \quad (2.3)$$

Note that “degree class \mathbf{k} ” means all the nodes with in-degree k_{in} and out-degree k_{out} ; $P(\mathbf{k})$ is the probability that a node is in degree class \mathbf{k} . Taking the average of Equation (2.2) for all nodes of the degree class \mathbf{k} , we obtain

$$\frac{1}{NP(\mathbf{k})} \sum_{i \in \mathbf{k}} p_n(i) = \frac{q}{N} + \frac{(1-q)}{NP(\mathbf{k})} \sum_{i \in \mathbf{k}} \sum_{j: k_{\text{out}}(j) \neq 0} \frac{a_{ji}}{k_{\text{out}}(j)} p_{n-1}(j). \quad (2.4)$$

From Equation (2.3) we see that the left-hand side of Equation (2.4) is $\bar{p}_n(\mathbf{k})$. On the right-hand side, we split the sum over j into two sums, one over all degree classes \mathbf{k}' and the other over all the nodes within each degree class \mathbf{k}' . We get

$$\bar{p}_n(\mathbf{k}) = \frac{q}{N} + \frac{(1-q)}{NP(\mathbf{k})} \sum_{\mathbf{k}'} \frac{1}{k'_{\text{out}}} \sum_{i \in \mathbf{k}} \sum_{j \in \mathbf{k}'} a_{ji} p_{n-1}(j). \quad (2.5)$$

At this point, we perform our mean field approximation [Binney et al. 92], which consists of substituting the PageRank of the predecessor neighbors of node i by its mean value, that is,

$$\begin{aligned} \sum_{i \in \mathbf{k}} \sum_{j \in \mathbf{k}'} a_{ji} p_{n-1}(j) &\simeq \bar{p}_{n-1}(\mathbf{k}') \sum_{i \in \mathbf{k}} \sum_{j \in \mathbf{k}'} a_{ji} \\ &= \bar{p}_{n-1}(\mathbf{k}') E_{\mathbf{k}' \rightarrow \mathbf{k}}, \end{aligned} \quad (2.6)$$

¹For example, on hypercubic lattices, the mean field solution for the Ising spin model becomes exact in four dimensions, when each spin has eight neighbors.

where $E_{\mathbf{k}' \rightarrow \mathbf{k}}$ is the total number of links pointing from nodes of degree \mathbf{k}' to nodes of degree \mathbf{k} . This matrix can also be rewritten as

$$\begin{aligned} E_{\mathbf{k}' \rightarrow \mathbf{k}} &= k_{\text{in}} P(\mathbf{k}) N \frac{E_{\mathbf{k}' \rightarrow \mathbf{k}}}{k_{\text{in}} P(\mathbf{k}) N} \\ &= k_{\text{in}} P(\mathbf{k}) N P_{\text{in}}(\mathbf{k}' | \mathbf{k}), \end{aligned} \quad (2.7)$$

where $P_{\text{in}}(\mathbf{k}' | \mathbf{k})$ is the probability that a predecessor of a node belonging to degree class \mathbf{k} belongs to degree class \mathbf{k}' . The conditional probability $P_{\text{in}}(\mathbf{k}' | \mathbf{k})$ incorporates the *degree-degree correlation*, i.e., the correlation between the degree of a node and that of its neighbors (see [Pastor-Satorras and Vespignani 04], pp. 243–245). Using Equations (2.6) and (2.7) in Equation (2.5), we finally obtain

$$\bar{p}_n(\mathbf{k}) = \frac{q}{N} + (1 - q) k_{\text{in}} \sum_{\mathbf{k}'} \frac{P_{\text{in}}(\mathbf{k}' | \mathbf{k})}{k'_{\text{out}}} \bar{p}_{n-1}(\mathbf{k}'), \quad (2.8)$$

which is a closed set of equations for the average PageRank of pages in the same degree class. When the network has degree-degree correlations, the solution of this equation is non-trivial, and the resulting PageRank can have a complex dependence on the degree. Direct measurements in our four Web samples, however, have revealed that degree-degree correlations are very weak in the Web [Serrano et al. 07]. In the particular case of networks without degree-degree correlations, the transition probability $P_{\text{in}}(\mathbf{k}' | \mathbf{k})$ becomes independent of \mathbf{k} and takes the simpler form

$$P_{\text{in}}(\mathbf{k}' | \mathbf{k}) = \frac{k'_{\text{out}} P(\mathbf{k}')}{\langle k_{\text{in}} \rangle},$$

where $\langle \cdot \rangle$ denotes the average value of the quantity in brackets. Using this expression in Equation (2.8) and taking the limit $n \rightarrow \infty$, we obtain

$$\bar{p}(\mathbf{k}) = \frac{q}{N} + \frac{1 - q}{N} \frac{k_{\text{in}}}{\langle k_{\text{in}} \rangle}, \quad (2.9)$$

that is, the average PageRank of nodes of degree class \mathbf{k} is independent of k_{out} and proportional to k_{in} .

2.2. Fluctuation Analysis

The formalism presented in the previous subsection gives a solution for the average PageRank of nodes of the same degree class but tells us nothing about how PageRank is distributed within one degree class. To fill this gap, we extend our mean field approach to the fluctuations within a degree class. To this end, we

first start by taking the square of Equation (2.2):

$$\begin{aligned}
 p_n^2(i) &= \frac{q^2}{N^2} + \frac{2q(1-q)}{N} \sum_j \frac{a_{ji}}{k_{\text{out}}(j)} p_{n-1}(j) \\
 &\quad + (1-q)^2 \sum_j \frac{a_{ji}}{k_{\text{out}}^2(j)} p_{n-1}^2(j) \\
 &\quad + (1-q)^2 \sum_{j \neq j'} \frac{a_{ji} a_{j'i}}{k_{\text{out}}(j) k_{\text{out}}(j')} p_{n-1}(j) p_{n-1}(j'). \quad (2.10)
 \end{aligned}$$

As in the previous calculation, we take the average over degree classes of the square of PageRank and define

$$\overline{p}_n^2(\mathbf{k}) \equiv \frac{1}{NP(\mathbf{k})} \sum_{i \in \mathbf{k}} p_n^2(i).$$

Taking this average in Equation (2.10) and rearranging terms, we get

$$\begin{aligned}
 \overline{p}_n^2(\mathbf{k}) &= \frac{q^2}{N^2} + \frac{2q(1-q)}{N} k_{\text{in}} \sum_{\mathbf{k}'} \frac{P_{\text{in}}(\mathbf{k}'|\mathbf{k})}{k'_{\text{out}}} \overline{p}_{n-1}(\mathbf{k}') \\
 &\quad + (1-q)^2 k_{\text{in}} \sum_{\mathbf{k}'} \frac{P_{\text{in}}(\mathbf{k}'|\mathbf{k})}{k'^2_{\text{out}}} \overline{p}_{n-1}^2(\mathbf{k}') \\
 &\quad + (1-q)^2 k_{\text{in}} (k_{\text{in}} - 1) \cdot \sum_{\mathbf{k}'} \sum_{\mathbf{k}''} \frac{P_{\text{in}}(\mathbf{k}', \mathbf{k}''|\mathbf{k})}{k'_{\text{out}} k''_{\text{out}}} \overline{p}_{n-1}(\mathbf{k}') \overline{p}_{n-1}(\mathbf{k}''),
 \end{aligned}$$

where we have again used the mean field approach. The probability $P_{\text{in}}(\mathbf{k}', \mathbf{k}''|\mathbf{k})$ is the joint probability that a node of degree \mathbf{k} has simultaneously one predecessor of degree \mathbf{k}' and another of degree \mathbf{k}'' . Let us make the reasonable independence assumption that this joint distribution factors as $P_{\text{in}}(\mathbf{k}', \mathbf{k}''|\mathbf{k}) = P_{\text{in}}(\mathbf{k}'|\mathbf{k}) P_{\text{in}}(\mathbf{k}''|\mathbf{k})$. In this case we can write an equation for the standard deviation within a degree class, $\sigma_n^2(\mathbf{k}) = \overline{p}_n^2(\mathbf{k}) - \overline{p}_n^2(\mathbf{k})$, as follows:

$$\begin{aligned}
 \frac{\sigma_n^2(\mathbf{k})}{(1-q)^2} &= k_{\text{in}} \sum_{\mathbf{k}'} \frac{P_{\text{in}}(\mathbf{k}'|\mathbf{k})}{k'^2_{\text{out}}} \sigma_{n-1}^2(\mathbf{k}') \\
 &\quad + k_{\text{in}} \sum_{\mathbf{k}'} \frac{P_{\text{in}}(\mathbf{k}'|\mathbf{k})}{k'^2_{\text{out}}} \overline{p}_{n-1}^2(\mathbf{k}') \\
 &\quad - k_{\text{in}} \left[\sum_{\mathbf{k}'} \frac{P_{\text{in}}(\mathbf{k}'|\mathbf{k})}{k'_{\text{out}}} \overline{p}_{n-1}(\mathbf{k}') \right]^2. \quad (2.11)
 \end{aligned}$$

In the case of uncorrelated networks, this equation can be analytically solved in the limit $n \rightarrow \infty$:

$$\sigma^2(\mathbf{k}) = \frac{(1-q)^2}{N^2 \langle k_{\text{in}} \rangle^2} \frac{1}{\langle k_{\text{in}} \rangle} \frac{\left\langle \frac{(q \langle k_{\text{in}} \rangle + (1-q) k_{\text{in}})^2}{k_{\text{out}}} \right\rangle - 1}{1 - \frac{(1-q)^2}{\langle k_{\text{in}} \rangle} \left\langle \frac{k_{\text{in}}}{k_{\text{out}}} \right\rangle} k_{\text{in}}. \quad (2.12)$$

In the case of the Web, the heavy tail of the in-degree distribution and the high average in-degree allow us to simplify this expression to

$$\sigma^2(\mathbf{k}) \simeq \frac{(1-q)^4}{N^2 \langle k_{\text{in}} \rangle^3} \left\langle \frac{k_{\text{in}}^2}{k_{\text{out}}} \right\rangle k_{\text{in}}.$$

For greater in-degrees, the coefficient of variation is

$$\frac{\sigma(\mathbf{k})}{\bar{p}(\mathbf{k})} \simeq (1-q) \left[\left\langle \frac{k_{\text{in}}^2}{k_{\text{out}}} \right\rangle \frac{1}{\langle k_{\text{in}} \rangle k_{\text{in}}} \right]^{1/2}. \quad (2.13)$$

The factor $\left\langle \frac{k_{\text{in}}^2}{k_{\text{out}}} \right\rangle$ in this expression can be large when the network has a long-tailed degree distribution, which implies that the relative fluctuations are large for lesser in-degrees. Therefore, the true PageRank of pages with lesser in-degree may differ significantly from its mean field approximation. For greater in-degrees, however, the relative fluctuations become less important—due to the factor k_{in} in the denominator—and the average PageRank from Equation (2.9) gives a good approximation. This can be seen by analyzing the coefficient of variation for the nodes with the maximum degree $k_{\text{in}}^{\text{max}}$. Assuming that k_{out} is weakly correlated with k_{in} , the coefficient $\left\langle \frac{k_{\text{in}}^2}{k_{\text{out}}} \right\rangle$ scales with the maximum in-degree as $(k_{\text{in}}^{\text{max}})^{3-\gamma}$ and the coefficient of variation as $(k_{\text{in}}^{\text{max}})^{1-\gamma/2}$. Since $\gamma > 2$, the relative fluctuations go to 0. Then, for lesser in-degrees, we expect PageRank to be distributed according to a power law; for intermediate in-degrees, according to a distribution peaked at the average mean field value plus a power law tail; and for greater in-degrees, according to a Gaussian distribution centered around the predicted mean field solution. Note that the expression in Equation (2.13) relates to the relative fluctuations within a degree class, rather than across the entire graph. Since PageRank is distributed according to a power law with γ close to 2, the overall fluctuations diverge in the limit of infinite graph size.

3. Results

For an empirical validation of the theoretical predictions in the previous section, we analyzed four samples of the Web graph. Two of them were obtained by crawls

Data Set	WB 2001	.uk 2002	WB 2003	.it 2004
Number of Pages	8.1×10^7	1.9×10^7	4.9×10^7	4.1×10^7
Number of Links	7.5×10^8	2.9×10^8	1.2×10^9	1.1×10^9
$\langle k \rangle$	9.34	15.78	24.05	27.50

Table 1. Number of pages, links, and average degree ($\langle k \rangle = \langle k_{in} \rangle = \langle k_{out} \rangle$) for the four data sets we have analyzed.

performed in 2001 and 2003 by the WebBase collaboration [Hirai et al. 00]. The other two were collected by the WebGraph project [Boldi and Vigna 04]; the pages belong to two national domains, .uk (2002) and .it (2004), respectively. In Table 1, we list the number of vertices and edges and the average degree for each data set.

We calculated PageRank with the standard iterative procedure; the factor q was set to 0.15, as in the original paper by Brin and Page [Brin and Page 98] and many successive studies. The convergence of the algorithm is very quick: In each case, fewer than 100 iterations were sufficient to determine the result with a relative accuracy of 10^{-5} for each vertex. (See [Berkhin 05] for an in-depth survey of PageRank computation.)

In Figure 1, we show the cumulative distributions of PageRank, i.e., the function $R(p)$ representing the probability that PageRank exceeds the value p . Using the cumulative distribution reduces the noise due to fluctuations at greater

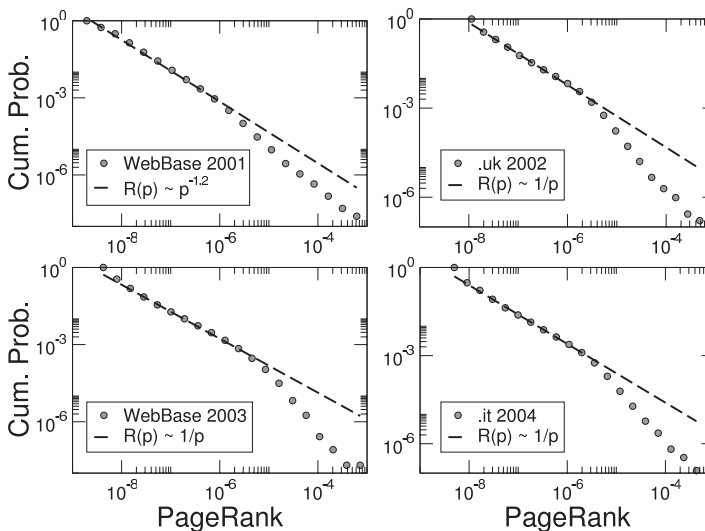


Figure 1. Cumulative distributions of PageRank.

Data Set	WB 2001	.uk 2002	WB 2003	.it 2004
γ	2.2 ± 0.1	2.0 ± 0.1	2.0 ± 0.1	2.0 ± 0.1
ρ	0.538	0.554	0.483	0.733

Table 2. Exponents of the power law part of the PageRank distribution and linear correlation coefficients between PageRank and in-degree.

PageRank values. In all four cases, we obtained a pattern with a broad tail. The initial part of the cumulative distribution fits well a power law $p^{-\beta}$ with exponent β between 1.0 and 1.2. The exponents for the actual PageRank distribution are $\gamma = \beta + 1$, so they range from 2.0 to 2.2, in agreement with other studies [Pandurangan et al. 02, Donato et al. 04]. The right-most part of each curve, corresponding to the pages with greatest PageRank, decreases faster. For the WebBase sample of 2001, the tail of the curve, up to the last point, fits well a power law with exponent $\gamma \approx 2.6$; in the other cases we see evidence of an exponential cutoff.

We also calculated the linear correlation coefficient between PageRank and in-degree. In Table 2, we list Pearson's ρ together with the slope of the power law portions of the PageRank distributions. The correlation between PageRank and in-degree is rather strong, in contrast to the findings of [Pandurangan et al. 02] and especially [Donato et al. 04], but in agreement with [Nakamura 03] and consistent with the high correlation observed between in-degree and Kleinberg's authority score [Donato et al. 05].

Let us now validate the expression derived from our mean field analysis for the average PageRank. We solved Equation (2.8) with an iterative procedure analogous to the one we used to calculate PageRank. We now look for the vector $\bar{p}(\mathbf{k})$, defined for all pairs $\mathbf{k} \equiv (k_{\text{in}}, k_{\text{out}})$ that occur in the network. Since PageRank is a probability, it must be normalized so that its sum over all vertices of the network is 1. So we initialized the vector with the constant $\bar{p}_0(\mathbf{k}) = 1/N$, and plugged it into the right-hand side of Equation (2.8) to get the first approximation, $\bar{p}_1(\mathbf{k})$. We then used $\bar{p}_1(\mathbf{k})$ as input to get $\bar{p}_2(\mathbf{k})$, and so on. We remark that the expression of the probability $P_{\text{in}}(\mathbf{k}'|\mathbf{k})$ is not a necessary ingredient of the calculation. In fact, the sum on the right-hand side of Equation (2.8) is just the average value of $\bar{p}_{n-1}(\mathbf{k}')/k'_{\text{out}}$ among all predecessors of vertices with degree \mathbf{k} . The algorithm leads to convergence within a few iterations (we never needed more than 20).

In Figure 2, we compare the values of $\bar{p}(\mathbf{k})$ calculated from Equation (2.8) with the corresponding empirical values. Here we averaged $\bar{p}(\mathbf{k})$ over out-degree, so it only depends on the in-degree k_{in} . The variation of $\bar{p}(\mathbf{k})$ with k_{out} (for fixed k_{in}) turns out to be very small. The scatter plots of Figure 2 show that




Figure 2. Scatter plots of the empirical average PageRank per degree class versus our mean field (MF) estimate.

the mean field approximation gives excellent results: the points are very tightly concentrated about each frame bisector, drawn as a guide to the eye. The mean field approximations of PageRank are well correlated with the degree-class averages of the actual PageRank. For example, Pearson's correlation coefficients are $\rho = 0.9$ and $\rho = 0.7$ for the WebBase crawls of 2001 and 2003, respectively.

Next let us analyze explicitly the relation between PageRank and in-degree. To plot the function $\bar{p}(k_{\text{in}})$ directly is not very helpful because the wide fluctuations of PageRank within each degree class would mystify the pattern for large values of k_{in} . So we average PageRank within bins of in-degree, which is the standard procedure to derive trends from scatter plots (see [Pastor-Satorras and Vespignani 04], pp. 240–242). As both PageRank and in-degree are power-law distributed, we use logarithmic bins; the multiplicative factor for the bin size is 1.3.

The resulting patterns for our four Web samples are presented in Figure 3. The empirical curves are rather smooth, and show that the average PageRank (per degree class) is an increasing function of in-degree. The relation between the two variables is approximately linear for greater in-degrees. This is exactly what we would expect if the degrees of pages were uncorrelated with those of their neighbors in the Web graph (see Section 2.1). In such a case, the relation between PageRank and in-degree is given by Equation (2.9). Indeed, the comparison of the empirical data with the curves of Equation (2.9) in Figure 3 is quite good for all data sets. This provides indirect confirmation that the Web graph is

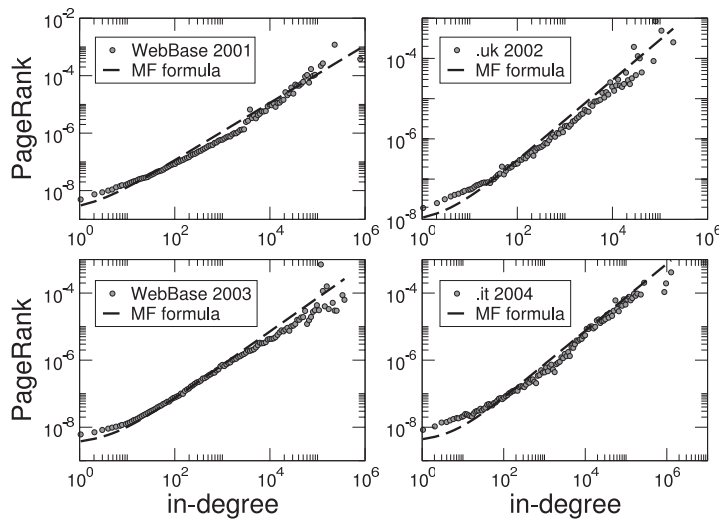


Figure 3. PageRank versus in-degree; the dashed line is the approximation given by the closed formula of Equation (2.9).

an essentially uncorrelated graph [Serrano et al. 07]. Most importantly, the average PageRank of a page with in-degree k_{in} is well approximated by the simple expression of Equation (2.9). The possible applications of this result are examined in the next section.

Let us analyze the empirical fluctuations of PageRank about its mean value. We anticipated in Section 2.2 that we expect large fluctuations for lesser values of k_{in} , due to the greater value of the second moment of the in-degree distribution, and that the relative size of the fluctuations should decrease as k_{in} increases (Equation (2.13)). Figure 4 confirms our prediction. We plotted the coefficient of variation $\sigma(k_{\text{in}})/\bar{p}(k_{\text{in}})$ as a function of k_{in} , once again averaging over out-degree. The trend is clearly decreasing as k_{in} increases. The fluctuations of the data points are due to degree-degree correlations (even if they are small, they are not completely negligible). We also derived mean field estimates for the coefficient of variation. Rather than solving the complete Equation (2.11), we used the coefficient of variation for an uncorrelated network, given by the ratio between $\sigma(k_{\text{in}})$ from Equation (2.12) and $\bar{p}(k_{\text{in}})$ from Equation (2.9). Nevertheless, the agreement between our approximated estimates and the empirical results in Figure 4 is very good except for high k_{in} , where we have an insufficient number of points in each degree class, leading to high fluctuations.

Finally, we test our prediction for the distribution of PageRank within a degree class. In Figure 5, we plot the PageRank distributions for four classes, corresponding to in-degree 1, 10, 100, and 1000. The data refer to the WebBase

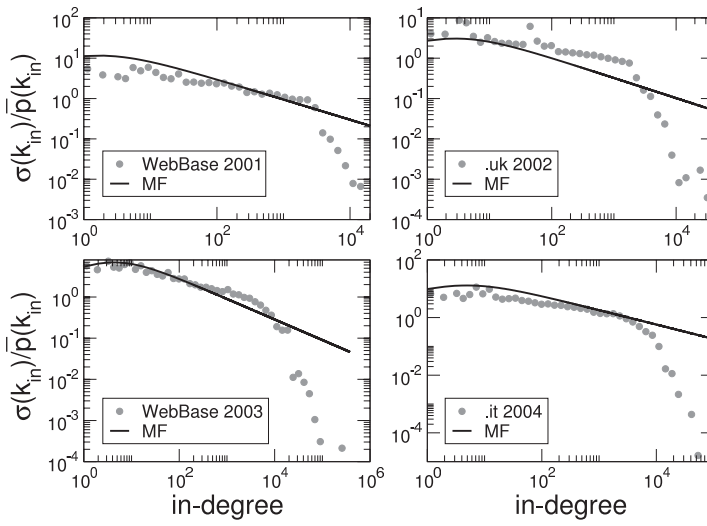


Figure 4. Coefficient of variation of PageRank versus in-degree.




Figure 5. Distributions of PageRank for four degree classes in the link graph from the WebBase 2003 crawl.

sample of 2003, but we found the same trend for the other three data sets. We see that, for low in-degrees ($k_{\text{in}} = 1, 10$), the distribution is a power law. The exponent is essentially the same (2.3 ± 0.1 , as for the other samples). For higher degrees ($k_{\text{in}} = 100, 1000$), however, the distribution changes from a power law to a hybrid distribution between a Gaussian and a power law. The Gaussian is signaled by the peak, which in the double logarithmic scale of the plot appears quite flat; the power law is manifest in the long tail of the distribution. The power law tails are all approximately parallel to each other, i.e., the exponent is the same for all curves.

4. Applications to the Live Web

Knowing the relationship between PageRank and in-degree has potential applications for the Web graph. It is vital for many service and information providers to have good rankings by major search engines for relevant keywords, given that search engines are an important way in which Internet users find and visit Web sites [RustyBrick 05, Qiu et al. 05]. This situation has made PageRank, used as a global ranking factor by Google, a valuable good controlled by a major search engine. Consequently a demand has emerged for companies that perform *search engine optimization* or *search engine marketing* on behalf of business clients. The goal is to increase the rankings of their pages, thus directing traffic to their sites [Sullivan 07].

Search engine marketers have partial knowledge of how search engines rank pages. They have access to undocumented tools to measure PageRank, such as the Google toolbar. Through experience and empirical tests, they can reverse-engineer some important ranking factors. After inspecting the hundreds of bulletin boards and blogs maintained by search engine marketers, however, it is evident that their work is guided largely by guesswork, trial, and error. Nevertheless search engine optimization has grown to be a healthy industry as shown in a recent study [iProspect 09]. Search engine marketing has even assumed ethical and legal ramifications, as a sort of arms race has ensued between marketers who want to increase their clients' rankings and search engines who want to maintain the integrity of their systems. The term *spamdexing* refers to those means of promoting Web sites that search engines deem unethical and worthy of blocking [Gori and Witten 05].

In the previous section, we showed that the average PageRank of a page with in-degree k_{in} can be well approximated by the closed formula in Equation (2.9). We have also found that PageRank fluctuations about the average become less important for greater k_{in} ; for great enough k_{in} the PageRank of a page with

k_{in} in-links only depends on k_{in} , and Equation (2.9) should give at least the correct order of magnitude for its PageRank. These results suggest that Web authors may use local in-degree information as a proxy for estimating the global PageRank of their sites.

To use Equation (2.9) for the Web we need to know the total number N of Web pages indexed by a search engine (say, Google) and their average degree $\langle k_{\text{in}} \rangle$. The size of the Google index was published until recently; we use the last reported number, $N \simeq 8.1 \times 10^9$. The average degree is not known; the best we can do is extract it from samples of the Web graph. Our data sets do not deliver a unique value for $\langle k_{\text{in}} \rangle$, but they agree on the order of magnitude (see Table 1). Hereafter we use $\langle k_{\text{in}} \rangle = 10$.

Let us now consider whether Equation (2.9) can be useful in the live Web. Ideally, we should compare the PageRank values of a list of Web pages with the corresponding values derived through our formula. Unfortunately, the real PageRank values calculated by a search engine such as Google are not accessible, so we need a different strategy. The simplest choice is to focus on rank rather than PageRank. We know that Google ranks Web pages according to their PageRank values as well as other features that do not depend on Web topology. The latter features are not disclosed; in the following we disregard them and assume for simplicity that the ranking of a Web page depends exclusively on its PageRank value. There is a simple relation between the PageRank p of a Web page and the rank R of that page. The Zipf function $R(p)$ is simply proportional to the cumulative distribution of PageRank. Since the PageRank distribution is approximately a power law with exponent $\gamma \simeq 2.1$ (see Section 3), we find that

$$R(p) \simeq A p^{-\beta}, \quad (4.1)$$

where $\beta = \gamma - 1 \simeq 1.1$ and A is a proportionality constant. Equation (4.1) can be empirically tested. Figure 6 shows the pattern for the WebBase sample of 2003. The approximation of Equation (4.1) (with $\beta = 1.1$) reproduces the data for over three orders of magnitude.

The rank R referred to above is the global rank of a page of PageRank p , i.e., its position in the list containing all pages of the Web in decreasing order of PageRank. More interesting for information providers and search engine marketers is the rank within hit lists returned for actual queries, where only a limited number of result pages appear. We need a criterion to pass from the global rank R to the rank r within a query's hit list. A page with global rank R could appear at any position $r = 1, 2, \dots, n$ in a list with n hits. In our framework, pages differ only by their PageRank values (or, equivalently, by their in-degrees), as we neglect lexical and other features. Therefore we can assume that each Web page has the same probability of appearing in a hit list. This is a strong assumption,




Figure 6. Dependence of the rank of a page on its PageRank value.

but even if it may fail to describe what happens at the level of an individual query, it is a fair approximation when one considers a large number of queries. Under this hypothesis, the probability distribution of the possible positions is a Poissonian, and the expected local rank r of a page with global rank R is given by the mean value:

$$r = R \frac{n}{N}. \quad (4.2)$$

Now it is possible to test the applicability of Equation (2.9) to the Web. We are able to estimate the rank of a Web page within a hit list if we know the number of in-links k_{in} of the page and the number n of hits in the list. The procedure consists of three simple steps:

1. From k_{in} we calculate the PageRank p of the page according to Equation (2.9).
2. From p we determine the global rank R according to Equation (4.1).
3. From R and n we derive the local rank r according to Equation (4.2).

The combination of the three steps leads to the following expression of the local rank r as a function of k_{in} and n :

$$r = \frac{A n}{\left(\frac{q}{N} + \frac{1-q}{N \langle k_{\text{in}} \rangle} k_{\text{in}}\right)^{1.1} N}. \quad (4.3)$$




Figure 7. Density map of the scatter plot between predicted rank r_{est} and actual rank r_{emp} for 65,207 queries. The fraction of points in each log-size bin is expressed by the color, also on a logarithmic scale. The diagonal guide to the eye is $r_{\text{est}} = r_{\text{emp}}$.

The natural way to derive the parameter A would be to perform a fit of the empirical relation between global rank and PageRank, as we did in Figure 6. The result should then be extrapolated to the full Web graph. As it turns out, the A -value derived in this way strongly depends on the sample of the Web, so that one could do no better than estimating the order of magnitude of A . On the other hand, A is a simple multiplicative constant, and its value has no effect on the dependence of the local rank r on the variables k_{in} and n . Therefore, we decided to consider it as a free parameter, whose value is to be determined by the comparison with empirical data.

For our analysis, we used a set of 65,207 actual queries from a September 2001 AltaVista log. We submitted each query to the Google Search API,² and picked at random one of the pages of the corresponding hit list. For each selected page, we stored its actual rank r_{emp} within the hit list, as well as its number k_{in} of in-links, which was again determined through Google.³ The number n of hits of the list was also stored. Google (like other search engines) never displays more than 1000 results per query, so we always have $r_{\text{emp}} \leq 1000$. From k_{in}

²<http://code.google.com/apis/>

³The in-degree data provided by search engines is only an estimate of the true number. First, a search engine can know only of links from pages that it has crawled and indexed. Second, for performance reasons, the algorithms counting inlinks use various unpublished approximations based on sampling.

and n we estimated the theoretical rank r_{est} by means of Equation (4.3), and compared it with its empirical counterpart r_{emp} . The comparison can be seen in the scatter plot of Figure 7. Given the large number of queries and the broad range of rank values, we visualize the density of points in logarithmic bins. The region with highest density is a stripe centered on the diagonal line $r_{\text{est}} = r_{\text{emp}}$ by a suitable choice of A ($A = 1.5 \times 10^{-4}$). We conclude that the rank derived through Equation (4.3) is, in most cases, close to the empirical one. We stress that this result is not trivial, because

1. Web pages are not ranked exclusively according to PageRank,
2. we are neglecting PageRank fluctuations,
3. pages do not have the same probability of being relevant with respect to a query.

Given such coarse approximations, we are satisfied with finding a good correlation at least between the orders of magnitude of the predicted and actual rank. Indeed, Pearson's correlation coefficient between the log-values of r_{est} and r_{emp} is $\rho = 0.8$, and a linear regression yields $p < 10^{-4}$.

5. Discussion

The present study motivates further inquiry. The mean field approach provides a simple functional relationship between average PageRank, in-degree, and degree-degree correlations. The price one pays for using such a simple approximation is the neglect of the significant fluctuations of PageRank values around the mean field average within a degree class. For the majority of pages, having moderate PageRank, fluctuations are more important; the in-degree being equal, fluctuations make the difference between being linked by “good” and “bad” pages. A path we intend to follow is understanding how two pages with the same in-degree end up with very different values of PageRank, and how pages with higher PageRank are positioned in the complex architecture of the Web graph.

The approach described here naturally lends itself to applications other than Web information retrieval, such as crawling [Cho et al. 98], traffic prediction [Fortunato et al. 06], spamdexing mitigation [Gyongyi et al. 04], and bibliometrics [Chen et al. 07]. For example, pages whose actual PageRank is much lower than its mean field approximation may flag link farms and be considered likely spam candidates. In bibliometrics, the quality of papers is commonly assessed via the number of citations they receive; it would be useful to be able to rank papers with the same number of citations through their PageRank values. A

characterization of papers leading to high PageRank fluctuations would be useful in this domain as well.

A promising way to study fluctuations for a moderate increase in complexity could be to use the definition of Equation (2.1), where the value of PageRank on the righthand side is substituted by the mean field approximation. Further work is needed in this direction.

In this paper, we have quantitatively explored two key assumptions around the current search status quo, namely that PageRank is very different from in-degree due to its global nature and that PageRank cannot be easily guessed or approximated without global knowledge of the Web graph. We have shown that due to the weak degree-degree correlations in the Web link graph, PageRank is strongly correlated with in-degree, and thus the two measures provide very similar information, especially for the most popular pages. Further, we have introduced a general mean field approximation of PageRank that, in the specific case of the Web, allows estimation of PageRank from only local knowledge of in-degree. We have further quantified the fluctuations of the approximation, gauging the reliability of the estimate. Finally, we have validated the approach with a simple procedure that predicts how actual Web pages are ranked by Google in response to actual queries, using only knowledge about in-degree and the number of query results.

Our method has immediate application for information providers. For instance, the association between rank and in-degree allows one to deduce how many in-links would be needed for a new page to achieve a given rank among all pages that deal with the same topic. This is an issue of crucial economic impact. All companies that advertise their products and services online want their homepages to be among the top-ranked sites in their business sector. Suppose that someone wants his or her homepage H to appear among the top n pages about topic T . Our recipe is extremely simple and cheap, requiring the submission of two queries to the search engine:

1. Submit a query to Google (or another search engine) about topic T .
2. Find the number k_n of in-links for the n th page of the resulting hit list.
3. H needs at least k_n in-links to appear among the first n hits for topic T .

Of course, there are limits to this approach; we do not claim that the lower bound of the number of in-links can be taken as a safe guide. Indeed, we have neglected important factors, such as the role of page content, in retrieving and ranking results and the fluctuations of the mean field approximation of PageRank. Besides, the necessity of limiting our empirical analysis to no more than 1000 hits

per query enhances the final correlation between in-degree and PageRank and represents a bias that deserves further investigations.

Notwithstanding the above caveats, our results indicate that at least the order of magnitude should be a reliable reference point. This may be all that is necessary—knowing the difference between the need for 1000 or 1,000,000 links can be a crucial asset in planning and budgeting a marketing campaign. Is word of mouth sufficient, or is advertising required? Our approach provides a tool to answer this kind of question. In making such a tool available to search engine marketers and information providers alike, we hope to create a more level playing field, so that not only large and powerful organizations but also small communities with little or no marketing budget can make informed decisions about the management of their Web presence.

Acknowledgments. We thank A. Vespignani and M. Serrano, for helpful discussions. We are grateful to Google for extensive use of its Web API, to the WebBase and WebGraph projects for their crawl data, and to AltaVista for use of their query logs. This work is funded in part by the Spanish government's DGES grant FIS2004-05923-CO2-02 and Generalitat de Catalunya grant SGR00889 to MB, by NSF Career award 0348940 to FM, and by the Indiana University School of Informatics. A preliminary version of this paper appeared in the Proceedings of the Fourth Workshop on Algorithms and Models for the Web-Graph (WAW2006).

References

- [Amento et al. 00] B. Amento, L. Terveen, and W. Hill. "Does 'Authority' Mean Quality? Predicting Expert Quality Ratings of Web Documents." In *Proc. 23rd ACM SIGIR Conf. on Research and Development in Information Retrieval*, pp. 296–303. New York: ACM Press, 2000.
- [Berkhin 05] P. Berkhin. "A Survey on PageRank Computing." *Internet Mathematics*, 2:1 (2005), 73–120.
- [Binney et al. 92] J. Binney, N. Dowrick, A. Fisher, and M. Newman. *The Theory of Critical Phenomena*. Oxford, UK: Oxford University Press, 1992.
- [Boldi and Vigna 04] P. Boldi and S. Vigna. "The WebGraph Framework I: Compression Techniques." In *Proc. 13th Intl. Conf. on World Wide Web*, pp. 595–601. New York: ACM Press, 2004.
- [Brin and Page 98] S. Brin and L. Page. "The Anatomy of a Large-Scale Hypertextual Web Search Engine." *Computer Networks* 30:1–7 (1998), 107–117.
- [Chen et al. 07] P. Chen, H. Xie, S. Maslov, and S. Redner. "Finding Scientific Gems with Google." *Journal of Informetrics* 1:1 (2007), 8–15.
- [Cho et al. 98] J. Cho, H. Garcia-Molina, and L. Page. "Efficient Crawling through URL Ordering." *Computer Networks* 30:1–7 (1998), 161–172.

- [Del Corso et al. 05] G. M. D. Corso, A. Gulli, and F. Romani. “Fast PageRank Computation via a Sparse Linear System.” *Internet Mathematics* 2:3 (2005), 251–273.
- [Donato et al. 04] D. Donato, L. Laura, S. Leonardi, and S. Millozzi. “Large Scale Properties of the Webgraph.” *European Physical Journal B* 38:2 (2004), 239–243.
- [Donato et al. 05] D. Donato, S. Leonardi, and P. Tsaparas. “Stability and Similarity of Link Analysis Ranking Algorithms.” In *Proc. of International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 717–729. Lecture Notes in Computer Science 3508. Berlin-Heidelberg: Springer-Verlag, 2005.
- [Fortunato et al. 06] S. Fortunato, A. Flammini, F. Menczer, and A. Vespignani. “Topical Interests and the Mitigation of Search Engine Bias.” *Proc. Natl. Acad. Sci. USA* 103:34 (2006) 12684–12689.
- [Gori and Witten 05] M. Gori and I. Witten. “The Bubble of Web Visibility.” *Communications of the ACM* 48:3 (2005), 115–117.
- [Gyongyi et al. 04] Z. Gyongyi, H. Garcia-Molina, and J. Pedersen. “Combating Web spam with TrustRank.” In *Proc. 30th International Conference on Very Large Data Bases (VLDB)*, pp. 576–587. VLDB Endowment, 2004.
- [Hirai et al. 00] J. Hirai, S. Raghavan, H. Garcia-Molina, and A. Paepcke. “WebBase: A Repository of Web Pages.” In *Proc. 9th International World Wide Web Conference*, pp. 277–293. Amsterdam: North-Holland Publishing Co., 2000.
- [iProspect 09] iProspect. “Search Engine Marketing Research Studies & White Papers.” Available online (<http://www.iprospect.com/about/searchenginemarketingwhitepapers.htm>), 2009.
- [Litvak et al. 07] N. Litvak, W. R. W. Scheinhardt, and Y. Volkovich. “In-Degree and PageRank of Web Pages: Why Do They Follow Similar Power Laws?” *Internet Mathematics* 4:2 (2007), 173–196.
- [Nakamura 03] I. Nakamura. “Characterization of Topological Structure on Complex Networks.” *Physical Review E* 68 (2003) 045104.
- [Pandurangan et al. 02] G. Pandurangan, P. Raghavan, and E. Upfal. “Using PageRank to Characterize Web Structure.” In *Proc. 8th Annual International Conference on Combinatorics and Computing (COCOON)*, pp. 330–339. Berlin-Heidelberg: Springer-Verlag, 2002.
- [Pastor-Satorras and Vespignani 04] R. Pastor-Satorras and A. Vespignani. *Evolution and Structure of the Internet*. Cambridge, UK: Cambridge University Press, 2004.
- [Qiu et al. 05] F. Qiu, Z. Liu, and J. Cho. “Analysis of User Web Traffic with a Focus on Search Activities.” In *Proc. International Workshop on the Web and Databases (WebDB)*, pp. 103–108. Available online (http://webdb2005.uhasselt.be/webdb05_e proceedings.pdf), 2005.
- [Richardson et al. 06] M. Richardson, A. Prakash, and E. Brill. “Beyond PageRank: Machine Learning for Static Ranking.” In *Proceedings of the 15th International Conference on World Wide Web*, pp. 707–715. New York: ACM Press, 2006.
- [Serrano et al. 07] M. A. Serrano, A. Maguitman, M. Boguna, S. Fortunato, and A. Vespignani. “Decoding the Structure of the WWW: A Comparative Analysis of Web Crawls.” *ACM Trans. Web* 1:2 (2007), 10.

- [Sullivan 07] D. Sullivan. “Intro to Search Engine Optimization.” Available online (<http://searchenginewatch.com/2167921>), 2007.
- [Sullivan 06] D. Sullivan. “Nielsen NetRatings Search Engine Ratings.” Available online (<http://searchenginewatch.com/2156451>), 2006.
- [Upstill et al. 03a] T. Upstill, N. Craswell, and D. Hawking. “Predicting Fame and Fortune: PageRank or Indegree?” In *Proc. 8th Australasian Document Computing Symposium (ADCS)*, pp. 31–40. Sydney: Australian Computer Society, 2003.
- [Upstill et al. 03b] T. Upstill, N. Craswell, and D. Hawking. “Query-Independent Evidence in Home Page Finding.” *ACM Trans. Information Systems* 21:3 (2003), 286–313.
- [RustyBrick 05] RustyBrick. “User Navigation Behavior to Effect Link Popularity.” Available online (<http://www.seroundtable.com/archives/001901.html>), 2005.

Santo Fortunato, Complex Networks Lagrange Laboratory, ISI Foundation, Viale S. Severo 65, 10133 Torino, Italy (fortunato@isi.it)

Marián Boguñá, Departament de Física Fonamental, Universitat de Barcelona, 08028 Barcelona, Spain (marian.boguna@ub.edu)

Alessandro Flammini, School of Informatics, Indiana University, 919 East Tenth St., Bloomington, IN 47408 (aflammin@indiana.edu)

Filippo Menczer, School of Informatics, Indiana University, 919 East Tenth St., Bloomington, IN 47408 (fil@indiana.edu)

Received September 21, 2007; accepted March 19, 2008.