

## 5.1 PageRank

We begin with a portion of the history of search engines, in order to motivate the definition of PageRank,<sup>2</sup> a tool for evaluating the importance of Web pages in a way that it is not easy to fool. We introduce the idea of “random surfers,” to explain why PageRank is effective. We then introduce the technique of “tax-ation” or recycling of random surfers, in order to avoid certain Web structures

---

<sup>1</sup>Link spammers sometimes try to make their unethicity less apparent by referring to what they do as “search-engine optimization.”

<sup>2</sup>The term PageRank comes from Larry Page, the inventor of the idea and a founder of Google.

that present problems for the simple version of PageRank.

### 5.1.1 Early Search Engines and Term Spam

There were many search engines before Google. Largely, they worked by crawling the Web and listing the *terms* (words or other strings of characters other than white space) found in each page, in an inverted index. An *inverted index* is a data structure that makes it easy, given a term, to find (pointers to) all the places where that term occurs.

When a *search query* (list of terms) was issued, the pages with those terms were extracted from the inverted index and ranked in a way that reflected the use of the terms within the page. Thus, presence of a term in a header of the page made the page more relevant than would the presence of the term in ordinary text, and large numbers of occurrences of the term would add to the assumed relevance of the page for the search query.

As people began to use search engines to find their way around the Web, unethical people saw the opportunity to fool search engines into leading people to their page. Thus, if you were selling shirts on the Web, all you cared about was that people would see your page, regardless of what they were looking for. Thus, you could add a term like “movie” to your page, and do it thousands of times, so a search engine would think you were a terribly important page about movies. When a user issued a search query with the term “movie,” the search engine would list your page first. To prevent the thousands of occurrences of “movie” from appearing on your page, you could give it the same color as the background. And if simply adding “movie” to your page didn’t do the trick, then you could go to the search engine, give it the query “movie,” and see what page *did* come back as the first choice. Then, copy that page into your own, again using the background color to make it invisible.

Techniques for fooling search engines into believing your page is about something it is not, are called *term spam*. The ability of term spammers to operate so easily rendered early search engines almost useless. To combat term spam, Google introduced two innovations:

1. PageRank was used to simulate where Web surfers, starting at a random page, would tend to congregate if they followed randomly chosen outlinks from the page at which they were currently located, and this process were allowed to iterate many times. Pages that would have a large number of surfers were considered more “important” than pages that would rarely be visited. Google prefers important pages to unimportant pages when deciding which pages to show first in response to a search query.
2. The content of a page was judged not only by the terms appearing on that page, but by the terms used in or near the links to that page. Note that while it is easy for a spammer to add false terms to a page they control, they cannot as easily get false terms added to the pages that link to their own page, if they do not control those pages.

### Simplified PageRank Doesn't Work

As we shall see, computing PageRank by simulating random surfers is a time-consuming process. One might think that simply counting the number of in-links for each page would be a good approximation to where random surfers would wind up. However, if that is all we did, then the hypothetical shirt-seller could simply create a “spam farm” of a million pages, each of which linked to his shirt page. Then, the shirt page looks very important indeed, and a search engine would be fooled.

These two techniques together make it very hard for the hypothetical shirt vendor to fool Google. While the shirt-seller can still add “movie” to his page, the fact that Google believed what other pages say about him, over what he says about himself would negate the use of false terms. The obvious countermeasure is for the shirt seller to create many pages of his own, and link to his shirt-selling page with a link that says “movie.” But those pages would not be given much importance by PageRank, since other pages would not link to them. The shirt-seller could create many links among his own pages, but none of these pages would get much importance according to the PageRank algorithm, and therefore, he still would not be able to fool Google into thinking his page was about movies.

It is reasonable to ask why simulation of random surfers should allow us to approximate the intuitive notion of the “importance” of pages. There are two related motivations that inspired this approach.

- Users of the Web “vote with their feet.” They tend to place links to pages they think are good or useful pages to look at, rather than bad or useless pages.
- The behavior of a random surfer indicates which pages users of the Web are likely to visit. Users are more likely to visit useful pages than useless pages.

But regardless of the reason, the PageRank measure has been proved empirically to work, and so we shall study in detail how it is computed.

#### 5.1.2 Definition of PageRank

PageRank is a function that assigns a real number to each page in the Web (or at least to that portion of the Web that has been crawled and its links discovered). The intent is that the higher the PageRank of a page, the more “important” it is. There is not one fixed algorithm for assignment of PageRank, and in fact variations on the basic idea can alter the relative PageRank of any two pages. We begin by defining the basic, idealized PageRank, and follow it

by modifications that are necessary for dealing with some real-world problems concerning the structure of the Web.

Think of the Web as a directed graph, where pages are the nodes, and there is an arc from page  $p_1$  to page  $p_2$  if there are one or more links from  $p_1$  to  $p_2$ . Figure 5.1 is an example of a tiny version of the Web, where there are only four pages. Page  $A$  has links to each of the other three pages; page  $B$  has links to  $A$  and  $D$  only; page  $C$  has a link only to  $A$ , and page  $D$  has links to  $B$  and  $C$  only.




Figure 5.1: A hypothetical example of the Web

Suppose a random surfer starts at page  $A$  in Fig. 5.1. There are links to  $B$ ,  $C$ , and  $D$ , so this surfer will next be at each of those pages with probability  $1/3$ , and has zero probability of being at  $A$ . A random surfer at  $B$  has, at the next step, probability  $1/2$  of being at  $A$ ,  $1/2$  of being at  $D$ , and 0 of being at  $B$  or  $C$ .

In general, we can define the *transition matrix of the Web* to describe what happens to random surfers after one step. This matrix  $M$  has  $n$  rows and columns, if there are  $n$  pages. The element  $m_{ij}$  in row  $i$  and column  $j$  has value  $1/k$  if page  $j$  has  $k$  arcs out, and one of them is to page  $i$ . Otherwise,  $m_{ij} = 0$ .

**Example 5.1:** The transition matrix for the Web of Fig. 5.1 is

$$M = \begin{bmatrix} 0 & 1/2 & 1 & 0 \\ 1/3 & 0 & 0 & 1/2 \\ 1/3 & 0 & 0 & 1/2 \\ 1/3 & 1/2 & 0 & 0 \end{bmatrix}$$

In this matrix, the order of the pages is the natural one,  $A$ ,  $B$ ,  $C$ , and  $D$ . Thus, the first column expresses the fact, already discussed, that a surfer at  $A$  has a  $1/3$  probability of next being at each of the other pages. The second column expresses the fact that a surfer at  $B$  has a  $1/2$  probability of being next at  $A$  and the same of being at  $D$ . The third column says a surfer at  $C$  is certain to be at  $A$  next. The last column says a surfer at  $D$  has a  $1/2$  probability of being next at  $B$  and the same at  $C$ .  $\square$

The probability distribution for the location of a random surfer can be described by a column vector whose  $j$ th component is the probability that the surfer is at page  $j$ . This probability is the (idealized) *PageRank* function.

Suppose we start a random surfer at any of the  $n$  pages of the Web with equal probability. Then the initial vector  $\mathbf{v}_0$  will have  $1/n$  for each component. If  $M$  is the transition matrix of the Web, then after one step, the distribution of the surfer will be  $M\mathbf{v}_0$ , after two steps it will be  $M(M\mathbf{v}_0) = M^2\mathbf{v}_0$ , and so on. In general, multiplying the initial vector  $\mathbf{v}_0$  by  $M$  a total of  $i$  times will give us the distribution of the surfer after  $i$  steps.

To see why multiplying a distribution vector  $\mathbf{v}$  by  $M$  gives the distribution  $\mathbf{x} = M\mathbf{v}$  at the next step, we reason as follows. The probability  $\mathbf{x}_i$  that a random surfer will be at node  $i$  at the next step, is  $\sum_j m_{ij}\mathbf{v}_j$ . Here,  $m_{ij}$  is the probability that a surfer at node  $j$  will move to node  $i$  at the next step (often 0 because there is no link from  $j$  to  $i$ ), and  $\mathbf{v}_j$  is the probability that the surfer was at node  $j$  at the previous step.

This sort of behavior is an example of the ancient theory of *Markov processes*. It is known that the distribution of the surfer approaches a limiting distribution  $\mathbf{v}$  that satisfies  $\mathbf{v} = M\mathbf{v}$ , provided two conditions are met:

1. The graph is *strongly connected*; that is, it is possible to get from any node to any other node.
2. There are no *dead ends*: nodes that have no arcs out.

Note that Fig. 5.1 satisfies both these conditions.

The limit is reached when multiplying the distribution by  $M$  another time does not change the distribution. In other terms, the limiting  $\mathbf{v}$  is an eigenvector of  $M$  (an *eigenvector* of a matrix  $M$  is a vector  $\mathbf{v}$  that satisfies  $\mathbf{v} = \lambda M\mathbf{v}$  for some constant *eigenvalue*  $\lambda$ ). In fact, because  $M$  is *stochastic*, meaning that its columns each add up to 1,  $\mathbf{v}$  is the *principal* eigenvector (its associated eigenvalue is the largest of all eigenvalues). Note also that, because  $M$  is stochastic, the eigenvalue associated with the principal eigenvector is 1.

The principal eigenvector of  $M$  tells us where the surfer is most likely to be after a long time. Recall that the intuition behind PageRank is that the more likely a surfer is to be at a page, the more important the page is. We can compute the principal eigenvector of  $M$  by starting with the initial vector  $\mathbf{v}_0$  and multiplying by  $M$  some number of times, until the vector we get shows little change at each round. In practice, for the Web itself, 50–75 iterations are sufficient to converge to within the error limits of double-precision arithmetic.

**Example 5.2:** Suppose we apply the process described above to the matrix  $M$  from Example 5.1. Since there are four nodes, the initial vector  $\mathbf{v}_0$  has four components, each  $1/4$ . The sequence of approximations to the limit that we

### Solving Linear Equations

If you look at the 4-node “Web” of Example 5.2, you might think that the way to solve the equation  $\mathbf{v} = M\mathbf{v}$  is by Gaussian elimination. Indeed, in that example, we argued what the limit would be essentially by doing so. However, in realistic examples, where there are tens or hundreds of billions of nodes, Gaussian elimination is not feasible. The reason is that Gaussian elimination takes time that is cubic in the number of equations. Thus, the only way to solve equations on this scale is to iterate as we have suggested. Even that iteration is quadratic at each round, but we can speed it up by taking advantage of the fact that the matrix  $M$  is very sparse; there are on average about ten links per page, i.e., ten nonzero entries per column.

Moreover, there is another difference between PageRank calculation and solving linear equations. The equation  $\mathbf{v} = M\mathbf{v}$  has an infinite number of solutions, since we can take any solution  $\mathbf{v}$ , multiply its components by any fixed constant  $c$ , and get another solution to the same equation. When we include the constraint that the sum of the components is 1, as we have done, then we get a unique solution.

get by multiplying at each step by  $M$  is:

$$\begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} \begin{bmatrix} 9/24 \\ 5/24 \\ 5/24 \\ 5/24 \end{bmatrix} \begin{bmatrix} 15/48 \\ 11/48 \\ 11/48 \\ 11/48 \end{bmatrix} \begin{bmatrix} 11/32 \\ 7/32 \\ 7/32 \\ 7/32 \end{bmatrix} \cdots \begin{bmatrix} 3/9 \\ 2/9 \\ 2/9 \\ 2/9 \end{bmatrix}$$

Notice that in this example, the probabilities for  $B$ ,  $C$ , and  $D$  remain the same. It is easy to see that  $B$  and  $C$  must always have the same values at any iteration, because their rows in  $M$  are identical. To show that their values are also the same as the value for  $D$ , an inductive proof works, and we leave it as an exercise. Given that the last three values of the limiting vector must be the same, it is easy to discover the limit of the above sequence. The first row of  $M$  tells us that the probability of  $A$  must be 3/2 the other probabilities, so the limit has the probability of  $A$  equal to 3/9, or 1/3, while the probability for the other three nodes is 2/9.

This difference in probability is not great. But in the real Web, with billions of nodes of greatly varying importance, the true probability of being at a node like `www.amazon.com` is orders of magnitude greater than the probability of typical nodes.  $\square$

### 5.1.3 Structure of the Web

It would be nice if the Web were strongly connected like Fig. 5.1. However, it is not, in practice. An early study of the Web found it to have the structure shown in Fig. 5.2. There was a large strongly connected component (SCC), but there were several other portions that were almost as large.

1. The *in-component*, consisting of pages that could reach the SCC by following links, but were not reachable from the SCC.
2. The *out-component*, consisting of pages reachable from the SCC but unable to reach the SCC.
3. *Tendrils*, which are of two types. Some tendrils consist of pages reachable from the in-component but not able to reach the in-component. The other tendrils can reach the out-component, but are not reachable from the out-component.




Figure 5.2: The “bowtie” picture of the Web

In addition, there were small numbers of pages found either in

- (a) *Tubes*, which are pages reachable from the in-component and able to reach the out-component, but unable to reach the SCC or be reached from the SCC.
- (b) Isolated components that are unreachable from the large components (the SCC, in- and out-components) and unable to reach those components.

Several of these structures violate the assumptions needed for the Markov-process iteration to converge to a limit. For example, when a random surfer enters the out-component, they can never leave. As a result, surfers starting in either the SCC or in-component are going to wind up in either the out-component or a tendril off the in-component. Thus, no page in the SCC or in-component winds up with any probability of a surfer being there. If we interpret this probability as measuring the importance of a page, then we conclude falsely that nothing in the SCC or in-component is of any importance.

As a result, PageRank is usually modified to prevent such anomalies. There are really two problems we need to avoid. First is the dead end, a page that has no links out. Surfers reaching such a page disappear, and the result is that in the limit no page that can reach a dead end can have any PageRank at all. The second problem is groups of pages that all have outlinks but they never link to any other pages. These structures are called *spider traps*.<sup>3</sup> Both these problems are solved by a method called “taxation,” where we assume a random surfer has a finite probability of leaving the Web at any step, and new surfers are started at each page. We shall illustrate this process as we study each of the two problem cases.

### 5.1.4 Avoiding Dead Ends

Recall that a page with no link out is called a dead end. If we allow dead ends, the transition matrix of the Web is no longer stochastic, since some of the columns will sum to 0 rather than 1. A matrix whose column sums are at most 1 is called *substochastic*. If we compute  $M^i \mathbf{v}$  for increasing powers of a substochastic matrix  $M$ , then some or all of the components of the vector go to 0. That is, importance “drains out” of the Web, and we get no information about the relative importance of pages.


**Example 5.3:** In Fig. 5.3 we have modified Fig. 5.1 by removing the arc from  $C$  to  $A$ . Thus,  $C$  becomes a dead end. In terms of random surfers, when a surfer reaches  $C$  they disappear at the next round. The matrix  $M$  that describes Fig. 5.3 is

$$M = \begin{bmatrix} 0 & 1/2 & 0 & 0 \\ 1/3 & 0 & 0 & 1/2 \\ 1/3 & 0 & 0 & 1/2 \\ 1/3 & 1/2 & 0 & 0 \end{bmatrix}$$

---

<sup>3</sup>They are so called because the programs that crawl the Web, recording pages and links, are often referred to as “spiders.” Once a spider enters a spider trap, it can never leave.



Figure 5.3:  $C$  is now a dead end

Note that it is substochastic, but not stochastic, because the sum of the third column, for  $C$ , is 0, not 1. Here is the sequence of vectors that result by starting with the vector with each component  $1/4$ , and repeatedly multiplying the vector by  $M$ :

$$\begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} \begin{bmatrix} 3/24 \\ 5/24 \\ 5/24 \\ 5/24 \end{bmatrix} \begin{bmatrix} 5/48 \\ 7/48 \\ 7/48 \\ 7/48 \end{bmatrix} \begin{bmatrix} 21/288 \\ 31/288 \\ 31/288 \\ 31/288 \end{bmatrix} \cdots \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

As we see, the probability of a surfer being anywhere goes to 0, as the number of steps increase.  $\square$

There are two approaches to dealing with dead ends.

1. We can drop the dead ends from the graph, and also drop their incoming arcs. Doing so may create more dead ends, which also have to be dropped, recursively. However, eventually we wind up with a strongly-connected component, none of whose nodes are dead ends. In terms of Fig. 5.2, recursive deletion of dead ends will remove parts of the out-component, tendrils, and tubes, but leave the SCC and the in-component, as well as parts of any small isolated components.<sup>4</sup>
2. We can modify the process by which random surfers are assumed to move about the Web. This method, which we refer to as “taxation,” also solves the problem of spider traps, so we shall defer it to Section 5.1.5.

If we use the first approach, recursive deletion of dead ends, then we solve the remaining graph  $G$  by whatever means are appropriate, including the taxation method if there might be spider traps in  $G$ . Then, we restore the graph, but keep the PageRank values for the nodes of  $G$ . Nodes not in  $G$ , but with predecessors

<sup>4</sup>You might suppose that the entire out-component and all the tendrils will be removed, but remember that they can have within them smaller strongly connected components, including spider traps, which cannot be deleted.

all in  $G$  can have their PageRank computed by summing, over all predecessors  $p$ , the PageRank of  $p$  divided by the number of successors of  $p$  in the full graph. Now there may be other nodes, not in  $G$ , that have the PageRank of all their predecessors computed. These may have their own PageRank computed by the same process. Eventually, all nodes outside  $G$  will have their PageRank computed; they can surely be computed in the order opposite to that in which they were deleted.




Figure 5.4: A graph with two levels of dead ends

**Example 5.4:** Figure 5.4 is a variation on Fig. 5.3, where we have introduced a successor  $E$  for  $C$ . But  $E$  is a dead end, and when we remove it, and the arc entering from  $C$ , we find that  $C$  is now a dead end. After removing  $C$ , no more nodes can be removed, since each of  $A$ ,  $B$ , and  $D$  have arcs leaving. The resulting graph is shown in Fig. 5.5.




Figure 5.5: The reduced graph with no dead ends

The matrix for the graph of Fig. 5.5 is

$$M = \begin{bmatrix} 0 & 1/2 & 0 \\ 1/2 & 0 & 1 \\ 1/2 & 1/2 & 0 \end{bmatrix}$$

The rows and columns correspond to  $A$ ,  $B$ , and  $D$  in that order. To get the PageRanks for this matrix, we start with a vector with all components equal to  $1/3$ , and repeatedly multiply by  $M$ . The sequence of vectors we get is

$$\begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix} \begin{bmatrix} 1/6 \\ 3/6 \\ 2/6 \end{bmatrix} \begin{bmatrix} 3/12 \\ 5/12 \\ 4/12 \end{bmatrix} \begin{bmatrix} 5/24 \\ 11/24 \\ 8/24 \end{bmatrix} \dots \begin{bmatrix} 2/9 \\ 4/9 \\ 3/9 \end{bmatrix}$$

We now know that the PageRank of  $A$  is  $2/9$ , the PageRank of  $B$  is  $4/9$ , and the PageRank of  $D$  is  $3/9$ . We still need to compute PageRanks for  $C$  and  $E$ , and we do so in the order opposite to that in which they were deleted. Since  $C$  was last to be deleted, we know all its predecessors have PageRanks computed. These predecessors are  $A$  and  $D$ . In Fig. 5.4,  $A$  has three successors, so it contributes  $1/3$  of its PageRank to  $C$ . Page  $D$  has two successors in Fig. 5.4, so it contributes half its PageRank to  $C$ . Thus, the PageRank of  $C$  is  $\frac{1}{3} \times \frac{2}{9} + \frac{1}{2} \times \frac{3}{9} = 13/54$ .

Now we can compute the PageRank for  $E$ . That node has only one predecessor,  $C$ , and  $C$  has only one successor. Thus, the PageRank of  $E$  is the same as that of  $C$ . Note that the sums of the PageRanks exceed 1, and they no longer represent the distribution of a random surfer. Yet they do represent decent estimates of the relative importance of the pages.  $\square$

### 5.1.5 Spider Traps and Taxation

As we mentioned, a spider trap is a set of nodes with no dead ends but no arcs out. These structures can appear intentionally or unintentionally on the Web, and they cause the PageRank calculation to place all the PageRank within the spider traps.

**Example 5.5:** Consider Fig. 5.6, which is Fig. 5.1 with the arc out of  $C$  changed to point to  $C$  itself. That change makes  $C$  a simple spider trap of one node. Note that in general spider traps can have many nodes, and as we shall see in Section 5.4, there are spider traps with millions of nodes that spammers construct intentionally.

The transition matrix for Fig. 5.6 is

$$M = \begin{bmatrix} 0 & 1/2 & 0 & 0 \\ 1/3 & 0 & 0 & 1/2 \\ 1/3 & 0 & 1 & 1/2 \\ 1/3 & 1/2 & 0 & 0 \end{bmatrix}$$




Figure 5.6: A graph with a one-node spider trap

If we perform the usual iteration to compute the PageRank of the nodes, we get

$$\begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} \begin{bmatrix} 3/24 \\ 5/24 \\ 11/24 \\ 5/24 \end{bmatrix} \begin{bmatrix} 5/48 \\ 7/48 \\ 29/48 \\ 7/48 \end{bmatrix} \begin{bmatrix} 21/288 \\ 31/288 \\ 205/288 \\ 31/288 \end{bmatrix} \cdots \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}$$

As predicted, all the PageRank is at  $C$ , since once there a random surfer can never leave.  $\square$

To avoid the problem illustrated by Example 5.5, we modify the calculation of PageRank by allowing each random surfer a small probability of *teleporting* to a random page, rather than following an out-link from their current page. The iterative step, where we compute a new vector estimate of PageRanks  $\mathbf{v}'$  from the current PageRank estimate  $\mathbf{v}$  and the transition matrix  $M$  is

$$\mathbf{v}' = \beta M\mathbf{v} + (1 - \beta)\mathbf{e}/n$$

where  $\beta$  is a chosen constant, usually in the range 0.8 to 0.9,  $\mathbf{e}$  is a vector of all 1's with the appropriate number of components, and  $n$  is the number of nodes in the Web graph. The term  $\beta M\mathbf{v}$  represents the case where, with probability  $\beta$ , the random surfer decides to follow an out-link from their present page. The term  $(1 - \beta)\mathbf{e}/n$  is a vector each of whose components has value  $(1 - \beta)/n$  and represents the introduction, with probability  $1 - \beta$ , of a new random surfer at a random page.

Note that if the graph has no dead ends, then the probability of introducing a new random surfer is exactly equal to the probability that the random surfer will decide *not* to follow a link from their current page. In this case, it is reasonable to visualize the surfer as deciding either to follow a link or teleport to a random page. However, if there are dead ends, then there is a third possibility, which is that the surfer goes nowhere. Since the term  $(1 - \beta)\mathbf{e}/n$  does not depend on

the sum of the components of the vector  $\mathbf{v}$ , there will always be some fraction of a surfer operating on the Web. That is, when there are dead ends, the sum of the components of  $\mathbf{v}$  may be less than 1, but it will never reach 0.

**Example 5.6:** Let us see how the new approach to computing PageRank fares on the graph of Fig. 5.6. We shall use  $\beta = 0.8$  in this example. Thus, the equation for the iteration becomes

$$\mathbf{v}' = \begin{bmatrix} 0 & 2/5 & 0 & 0 \\ 4/15 & 0 & 0 & 2/5 \\ 4/15 & 0 & 4/5 & 2/5 \\ 4/15 & 2/5 & 0 & 0 \end{bmatrix} \mathbf{v} + \begin{bmatrix} 1/20 \\ 1/20 \\ 1/20 \\ 1/20 \end{bmatrix}$$

Notice that we have incorporated the factor  $\beta$  into  $M$  by multiplying each of its elements by  $4/5$ . The components of the vector  $(1 - \beta)\mathbf{e}/n$  are each  $1/20$ , since  $1 - \beta = 1/5$  and  $n = 4$ . Here are the first few iterations:

$$\begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} \begin{bmatrix} 9/60 \\ 13/60 \\ 25/60 \\ 13/60 \end{bmatrix} \begin{bmatrix} 41/300 \\ 53/300 \\ 153/300 \\ 53/300 \end{bmatrix} \begin{bmatrix} 543/4500 \\ 707/4500 \\ 2543/4500 \\ 707/4500 \end{bmatrix} \cdots \begin{bmatrix} 15/148 \\ 19/148 \\ 95/148 \\ 19/148 \end{bmatrix}$$

By being a spider trap,  $C$  has managed to get more than half of the PageRank for itself. However, the effect has been limited, and each of the nodes gets some of the PageRank.  $\square$

### 5.1.6 Using PageRank in a Search Engine

Having seen how to calculate the PageRank vector for the portion of the Web that a search engine has crawled, we should examine how this information is used. Each search engine has a secret formula that decides the order in which to show pages to the user in response to a search query consisting of one or more search terms (words). Google is said to use over 250 different properties of pages, from which a linear order of pages is decided.

First, in order to be considered for the ranking at all, a page has to have at least one of the search terms in the query. Normally, the weighting of properties is such that unless all the search terms are present, a page has very little chance of being in the top ten that are normally shown first to the user. Among the qualified pages, a score is computed for each, and an important component of this score is the PageRank of the page. Other components include the presence or absence of search terms in prominent places, such as headers or the links to the page itself.

### 5.1.7 Exercises for Section 5.1

**Exercise 5.1.1:** Compute the PageRank of each page in Fig. 5.7, assuming no taxation.




Figure 5.7: An example graph for exercises

**Exercise 5.1.2:** Compute the PageRank of each page in Fig. 5.7, assuming  $\beta = 0.8$ .

**! Exercise 5.1.3:** Suppose the Web consists of a *clique* (set of nodes with all possible arcs from one to another) of  $n$  nodes and a single additional node that is the successor of each of the  $n$  nodes in the clique. Figure 5.8 shows this graph for the case  $n = 4$ . Determine the PageRank of each page, as a function of  $n$  and  $\beta$ .




Figure 5.8: Example of graphs discussed in Exercise 5.1.3

**!! Exercise 5.1.4:** Construct, for any integer  $n$ , a Web such that, depending on  $\beta$ , any of the  $n$  nodes can have the highest PageRank among those  $n$ . It is allowed for there to be other nodes in the Web besides these  $n$ .

**! Exercise 5.1.5:** Show by induction on  $n$  that if the second, third, and fourth components of a vector  $\mathbf{v}$  are equal, and  $M$  is the transition matrix of Example 5.1, then the second, third, and fourth components are also equal in  $M^n \mathbf{v}$  for any  $n \geq 0$ .




Figure 5.9: A chain of dead ends

**Exercise 5.1.6:** Suppose we recursively eliminate dead ends from the graph, solve the remaining graph, and estimate the PageRank for the dead-end pages as described in Section 5.1.4. Suppose the graph is a chain of dead ends, headed by a node with a self-loop, as suggested in Fig. 5.9. What would be the PageRank assigned to each of the nodes?

**Exercise 5.1.7:** Repeat Exercise 5.1.6 for the tree of dead ends suggested by Fig. 5.10. That is, there is a single node with a self-loop, which is also the root of a complete binary tree of  $n$  levels.




Figure 5.10: A tree of dead ends