Optimal Link Bombs are Uncoordinated *

Sibel Adalı, Tina Liu, Malik Magdon-Ismail
Department of Computer Science, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180.
Email: {sibel, liut2, magdon}@cs.rpi.edu

Abstract

We analyze the recent phenomenon termed a *Link Bomb*, and investigate the optimal attack pattern for a group of web pages attempting to link bomb a specific web page. The typical *modus operandi* of a link bomb is to associate a particular page with a search text and then boost that page's *pagerank*. (The attacking pages can only control their own content and outgoing links.) Thus, when a search is initiated with the text, a high prominence will be given to the attacked page. We show that the best organization of links among the attacking group to maximize the increase in rank of the attacked node is the *direct individual attack*, where every attacker points directly to the victim and nowhere else. We also discuss optimal attack patterns for a group that wants to hide itself by not pointing directly to the victim. We quantify our results with experiments on a variety of random graph models.

1 Introduction

Generally, a search on a particular topic on a particular search engine (such as *Google*) will output a ranked list of relevent web pages. The prominence of a page in this listing is an important indicator of how many people will visit the page. For a commercial web site, its prominence with respect to product searches has important financial consequences, as does the prominence of a competitor's website with respect to slander about products. Prominence in rankings is prestigious and can add credibility [10].

As a result of the importance attached to one's pagerank, especially one's Google pagerank, artificial methods for boosting one's pagerank with respect to a particular topic have become an active area for discussion. A prominent case was an attempt to "bring down the White house" by giving high prominence (in fact the primo ranking) to a web-biography of the U.S. President with respect to the text "miserable failure" [10, 14]. Such attacks are generally termed *Google bombs* (named after the success of such attacks on Google rankings), which are attempts to give prominence to a particular web page with respect to a particular (usually derogatory) piece of text[1]. We study *link bombs* which attempt to alter the rankings obtained through the PageRank algorithm by manipulating the links - such attacks would be relevant to any search engine that uses such a page ranking algorithm. Link bombs need not be derogatory, for example, a web-retailer could also make use of link

*This work was partially supported by the National Science Foundation under grant EIA-0091505

[1] After the first attack (which was with respect to the text "talentless hack"), several other attacks also succeeded in raising the ranks of web pages with respect to specific keyword(s), in some cases using as few as 25 links. It has been argued that several factors contribute to the success of an attack, eg. the number and prominence of the attacking pages; the (un)popularity of the keyword. Many of these attacks were usually initiated by Blogs which tend to be updated often and have a lot of content. It has been argued that these factors contribute to the high prominence of Blogs which in turn have higher influence in the pagerank of other pages. Similarly, some of the keywords chosen were very rare in the web pages, such as "French Military Victories". However, even attacks using keywords as popular as "Weapons of Mass Destruction" have been successful (BBC News, Sunday, 7 December, 2003).

bombs to improve the prominence of its own website with respect to a particular topic(s). The link bombers are usually some (coordinated) set of web pages which add outgoing links to their web page. Some of these links will point to the attacked page, and contain the text they (the bombers) are trying to associate with the attacked page. The issue we address is how these bombers should organize their outgoing links in order to maximize the success of their link bomb.

There is currently a great deal of discussion on whether a link bomb can be considered an "undesirable" attack [14] that exploits a weakness in the pagerank algorithm [6, 12]. The pagerank algorithm assigns you a pagerank by considering the number and importance (according to PageRank) of web pages that point to you. Given that a search engine like Google currently ranks about 8 billion pages, one would expect that a very small number of web pages should not be able to change the ranking of a page dramatically, contrary to what has been observed. Thus, one motivation for studying the optimal attack is to determine specific abnormal but effective attack patterns that could be identified as artificial Google bombs.

We present results on the optimal link bomb. Specifically, the *attackers* are a set of web pages whose outgoing links can be manipulated, and the *victim* is the target web page to be bombed. Our main result is to establish the following theorem as a starting point for a discussion of accountability on linked structures such as the WWW,

Theorem. *The attack which maximizes the* **rank** *of the victim with respect to page rank is the direct individual attack.*

The *direct individual attack* is the attack in which every attacker points *only* to the victim and to no other page. In particular, in the optimal attack, none of the attackers point to each other. Thus, the optimal attack masquerades as a set of uncoordinated "random" nodes, all pointing to the same page. We also discuss optimal "disguised" attack patterns, in which none of the attackers wish to directly point to the victim – all paths from the attackers to the victim must be of at least some minimum length. In this case the optimal attack is still a direct individual attack, however now the attackers point to some other *intermediate node* (not the victim).

While the optimal attack is always the direct individual attack, the amount by which the direct individual attack surpasses other (more coordinated) attack patterns may depend on the nature of the graph. We give experimental results that quantify this phenomenon for a variety of different attack patterns. On certain random graph models of the Web, some coordinated attack patterns are almost as good as the direct individual attack, and can hence be used in place of the direct individual attack as a means of disguising the attack. While the effect of graph structure on the pagerank has been investigated in the literature [11, 6], to our knowledge, these are the first results regarding the effect of the graph structure on the effectiveness of link bombs.

Our results raise interesting questions such as how to detect and respond to link bomb attacks (in general this problem is NP-hard, see for example [16]). Since the attackers will have no visible associations amongst themselves, it is hard to detect and prove that they are participating in an attack. If the optimal attack were a tree structure, there would be a small set of nodes with high prominence that one might argue are "responsible" for the attack. The other nodes pointing to these nodes could also be held accountable aiding and abbeting the actions of the responsible nodes. Such accountability is not possible in an individual attack.

We proceed by first discussing some preliminary definitions, followed by a preview of our result for an isolated graph, in which the only nodes are the attackers and the victim. We then discuss general graphs, followed by some experimental results on a variety of random graph models. We conclude with a discussion of the implications of our results. (We omit technical proofs which can be found in a full version of this paper [1].)

2 Preliminaries

A search query on a set of keywords results in an ordered list of web pages $\mathcal{W} = \{\omega_i\}$. Each web page $\omega \in \mathcal{W}$ contains some or all of the keywords either in its text or in the text of a link that points from some other web page to ω. A scoring function is used to order the pages in $\mathcal{W}$. The most prominent page (page with the highest score) is given rank 1, etc.

Google [3] considers many factors in its scoring function, including: keyword frequency; relative locations of the keywords; the position and style of the keywords. An important factor in the scoring function is the *pagerank* which depends on how the web page is embedded in the entire graph of web pages. An early paper on the Google system [3] suggests that no one factor dominates the scoring function, however, the pagerank plays an important role. In this paper, we will concentrate only on the pagerank factor and discuss how it can be manipulated.

The *web graph* is a directed graph $G = (V, E)$ that models the World Wide Web. The vertex set V represents the pages and documents, and the edge set E represents the links between the pages and documents[2]. The edges are directed: if $(v_1, v_2) \in E$, then v_1 contains a link to v_2. In a web graph, the in-degree $indeg(v)$ of page v is the number of links that point to v and the out-degree $outdeg(v)$ is the number of links originating from v that point to other pages. A (directed) *path* of length ℓ is a sequence of vertices $v_0, v_1, \ldots, v_\ell$ with $(v_{i-1}, v_i) \in E$ for $i = 1, \ldots, \ell$. v_ℓ is the terminal node in the path, and $v_1, \ldots, v_{\ell-1}$ are intermediate nodes. We allow parallel edges between two vertices, but no self-loops.

The pagerank p_i models the probability that node i will be visited either by randomly navigating down links in the web graph or by randomly jumping to page i. Let α be the probability to navigate, and $1 - \alpha$ the probability to jump. Then the pageranks $\{p_j\}$ of the nodes in a graph simultaneously satisfy the set of linear equations[3]

$$p_i = \alpha \sum_{(v_j, v_i) \in E} \frac{p_j}{outdeg(v_j)} + \frac{1-\alpha}{N}. \qquad (1)$$

($0 \le \alpha \le 1$ and $N = |V|$.) The first term represents the probability to reach i by random navigation. An edge may appear multiple times if there are parallel links. The second term represents the probability to reach i by randomly jumping. Typically, $\alpha \in [0.85, 0.95]$. p_i is larger if v_i has a large in-degree, and its incoming links are from high pagerank nodes with small out-degree. The PageRank algorithm [12] is an iterative approach to solving these equations. The pageranks are all initialized to $p_i^0 = \frac{1}{N}$. The PageRank iteration [12] is given by

$$p_i^{t+1} = \alpha \sum_{(v_j, v_i) \in E} \frac{p_j^t}{outdeg(v_j)} + \frac{1-\alpha}{N}. \qquad (2)$$

p_i^t converges to the (unique) solution of (1). Every page can manipulate its outgoing links, but it cannot change its incoming links.

A *link bomb*, or *attack* occurs when a group of *attackers* $A = \{v_1, \ldots, v_K\}$ alters their outgoing links so as to boost the pagerank of a *victim* $v_0 \notin A$. Before the attack, if the edge set is E, then after the attack the edge set will be $\bar{E}$ where the only edges added or removed from E are of the form (v_i, u) where $1 \le i \le K$ and $u \in V$, i.e., the attackers may remove and/or add outgoing links only. After the attack, the new web graph is $\bar{G} = (V, \bar{E})$. Let p_i denote the pageranks in the original graph G (before the attack), and $\bar{p}_i$ the pageranks in $\bar{G}$ (after the attack). The *magnitude* of the attack $\Delta p_0 = \bar{p}_0 - p_0$ is the amount by which the pagerank of the victim increased, and is a measure of the success of the attack. In our analysis, we only consider the magnitude of the attack,

[2]Note that the definition of an edge is traditionally given by hyperlinks in a web page. However, it is also possible to count URLs in the body of a web page as links. The definition of what constitutes a link is usually application dependent.

[3]An alternative and common formulation of the pageranks in the literature is as the stationary distribution of a suitably defined finite irreducible Markov chain with transition matrix $P = (1-\alpha)M + \alpha U$, where U is a matrix of 1's. Many of our results could be obtained by analyzing how the stationary distribution changes under perturbations of P. Our approach is more graph theoretic, treating the problem as a flow.

and assume that all other factors entering into the scoring function are unchanged.

3 The Optimal Link Bomb

In this section, we investigate how to maximize the magnitude of the attack. In particular, we show that the effectiveness of the attack *does not* increase if the attackers try to coordinate the attack in some way, by introducing links among themselves in order to increase their ranks. (Recall that, incoming links from higher ranked pages are more beneficial to your rank.) First, we consider a simplified case, in which the attackers and the victim are isolated from the rest of the graph. We then consider the general case.

3.1 Isolated Graphs

We can restrict our attention to the vertex set composed of the attackers and the victim, $V = A \cup v_0$ (i.e., $N = |V| = K + 1$). Assume (for simplicity) that v_0 does not point to any member of A. We first consider some examples of attacks, before giving the general result. In all cases, all the attackers A point to the victim v_0, and what differentiates the attacks is how the attackers are themselves organized.

Direct Individual: The only links are to v_0.

Tree: The attackers form a tree. For any graph with a topological order, one can compute the page ranks efficiently (in linear time). For analysis purposes, we will specialize to a *star attack* in which $v_2 \ldots, v_K$ point to v_1 and all attackers point to v_0.

Cycle: The attackers form a cycle.

Complete: The attackers a complete graph.

By solving the linear system (1) for the graph resulting from each of these attacks, we obtain

Lemma 1 *For the isolated graph,*

$$\begin{aligned}
\bar{p}_0(\textit{individual}) &= p_0(1+\alpha K), \\
\bar{p}_0(\textit{star}) &= p_0\left(1+\tfrac{\alpha}{2}(K(1+\alpha)+1-\alpha)\right), \\
\bar{p}_0(\textit{cycle}) &= p_0\left(1+\tfrac{\alpha K}{2-\alpha}\right), \\
\bar{p}_0(\textit{complete}) &= p_0\left(1+\tfrac{\alpha K}{K(1-\alpha)+\alpha}\right),
\end{aligned}$$

where $p_0 = (1-\alpha)/(K+1)$.

Since $0 \le \alpha \le 1$, after some algebra, we obtain

Theorem 1 *For the isolated graph,*

$$\bar{p}_0(\textit{individual}) \ge \bar{p}_0(\textit{star}) \ge \bar{p}_0(\textit{cycle}) \ge \bar{p}_0(\textit{complete}).$$

In fact, the direct individual attack is optimal for the isolated graph:

Theorem 2 *For an isolated graph, p_0 is maximized (uniquely) by the individual attack.*

3.2 Arbitrary Graphs

When $v_0, \ldots, v_K$ are embedded in a larger graph G, the direct individual attack is still optimal. Intuitively, one can view the PageRank iteration (2) as sending a flow of pagerank down the directed edges. The maximum flow from v_i to v_0 occurs when v_i points directly to v_0, and to no other node – any other links divert the flow and lead to a lower magnitude attack. The following results will make this intuition more formal. We will generally refer to nodes which are neither the attackers nor the victim by w_j, and u_j will be used to refer to any node. The 1-neighborhood $N_1(v)$ of a node v is the set of nodes to which v points. $N_k(v)$ $(k > 1)$ is the set of k-neighborhood nodes: $u \in N_k(v)$ *iff* for some $w \in N_{k-1}(v)$, $(w,u) \in E$. Note that v could be in its own k-neighborhood for $k > 1$, and $N_0(v) = \{v\}$.

Consider attacker v_i, and, without loss of generality, assume it initially has no outgoing links. Suppose now that it adds δ outgoing edges. This results in $\frac{\alpha}{\delta}$ of its rank "flowing" along each of its edges to its

neighbors (note there may be parallel links). Thus, the rank increase for a 1-neighbor u_j is given by

$$\Delta_j^1 = \alpha \sum_{(v_i,u_j)\in E} \frac{p_i}{outdeg(v_i)},$$

where the superscript 1 indicates that u_j is a 1-neighbor, and j is an index that enumerates the 1-neighbors. The sum is over all parallel edges that v_i may have to u_j. This increase in rank in turn propagates to 2-neighbors, resulting in an increase in the rank of a 2-neighbor u_k by an amount

$$\Delta_k^2 = \alpha \sum_{\substack{(u_j,u_k)\in E \\ s.t\ u_j\in N_1(v_i)}} \frac{\Delta_j^1}{outdeg(u_j)}.$$

The sum is over all 1-neighbors pointing to u_k (including parallel edges). If the newly added edges create a path from v_i to v_0, then some amount of v_i's pagerank will propagate to v_0. We define Δ_j^l to be the change in the page rank of u_j from flow down all paths of length l from v_i to u_j,

$$\Delta_j^l = \alpha \sum_{\substack{(u_k,u_j)\in E \\ s.t\ u_k\in N_{l-1}(v_i)}} \frac{\Delta_k^{l-1}}{outdeg(u_k)}$$

Let $\delta(l)$ be total increase in page rank through paths of length l, $\delta(l) = \sum_j \Delta_j^l$. Since the pagerank increase attenuates by a factor α with each edge, we have the following lemma.

Lemma 2 *$\delta(l) \le \alpha^l p_i$, with equality* iff *$\delta(l-1) = \alpha^{l-1} p_i$ and for every $u_k \in N_{l-1}(v_i)$, $outdeg(u_k) > 0$.*

Let $\mathcal{S}$ be a set of nodes. A path q passes through $\mathcal{S}$ if some node of $\mathcal{S}$ is an intermediate node of q. A set of paths P pass through $\mathcal{S}$ if every path in P passes through $\mathcal{S}$. Let P_t be a collection of paths that passes through $\mathcal{S}$, with every path in P_t having the same terminal node $t \neq v_i$ (t is not an intermediate node of any path in P_t). We call t a *progeny* of $\mathcal{S}$ with respect to the paths P_t. Since every path passes through $\mathcal{S}$, some prefix of every path in P_t has a terminal node in $\mathcal{S}$. For each path $q \in P_t$, let $q_{\mathcal{S}}$ be a (any) prefix with terminal node in $\mathcal{S}$, and let $P_t(\mathcal{S})$ denote the collection of such distinct prefixes $\{q_{\mathcal{S}}\}$.

The *influence* $I(\mathcal{S}|P_t(\mathcal{S}))$ of v_i on $\mathcal{S}$ is the total flow of pagerank (summed over all nodes in $\mathcal{S}$) from v_i to $\mathcal{S}$ along the paths in $P_t(\mathcal{S})$ (which are (distinct) prefixes in P_t). The influence $I(t|P_t)$ of v_i on t is the total flow of pagerank that flows to t along the paths in P_t (which pass through $\mathcal{S}$). Every path in P_t has at least one additional edge compared with its corresponding prefix that terminates in $\mathcal{S}$, so the influence that propagates to t along P_t can be at most the influence that propagates to $\mathcal{S}$ along the paths in $P_t(\mathcal{S})$, attenuated by a factor α. We have the following lemma.

Lemma 3 *$I(t|P_t) \le \alpha I(\mathcal{S}|P_t(\mathcal{S}))$, independent of which prefixes are used in the construction of $P_t(\mathcal{S})$.*

We now consider v_i's attack on v_0. Let P denote the collection of all (distinct) paths from v_i to v_0 in which v_0 appears only as the terminal node, i.e., v_0 is not an intermediate node of any path in P. Note that if there are cycles in the graph, then P may contain an infinite number of paths. Let the flow of pagerank from v_i to v_0 down the paths in P be denoted Δ. There may be cycles containing v_0, in which case, the pagerank increase Δ will continue to flow around these cycles, back to v_0 increasing the pagerank further, i.e., Δ will be amplified by the cycles. Let Δp_0^i be v_i's contribution to the magnitude of the attack,

$$\Delta p_0^i(\Delta) = \Delta + \mathrm{amp}(\Delta),$$

where $\mathrm{amp}(\Delta)$ is the amplification due to the cycles that contain v_0. The larger Δ, the larger will be the amplification of Δ,

Lemma 4 *$\Delta p_0^i(\Delta)$ is monotonically increasing.*

Lemmas 2, 3 and 4 are the main tools we will need to prove our main result, namely that the individual attack is optimal. By Lemma 4, since Δp_0^i is monotonically increasing in Δ, Δp_0^i will be maximized when Δ is maximized. Δ is given by the sum of the flows of pagerank from v_i to v_0 along the paths in P, therefore we only need to consider this flow.

Let ℓ be the length of the shortest path in P (there may be many such shortest paths). Consider the set L of all distinct paths of length ℓ originating at v_i. Some of these paths have terminal node v_0. We now restrict our attention to the set L' containing those paths in L which do not have terminal node v_0. Note that none of the paths in L' can have v_0 as an intermediate node since the shortest path from v_i to v_0 has length ℓ. Let Let $\mathcal{S}$ denote the set of terminal nodes in L'. Partition P into two disjoint sets, P_ℓ and $P_{>\ell}$, where P_ℓ contains the paths in P with length ℓ and $P_{>\ell}$ the paths with length $> \ell$. Every path in $P_{>\ell}$ must pass through at least one of the nodes in $\mathcal{S}$, therefore $P_{>\ell}$ passes through $\mathcal{S}$. Every path in $P_{>\ell}$ has terminal node v_0, and v_0 does not appear as an intermediate node in any of these paths. Thus, v_0 is a progeny of $\mathcal{S}$ with respect to $P_{>\ell}$. Every path in $P_{>\ell}$ has a prefix of length ℓ with terminal node in $\mathcal{S}$. Collect these distinct prefixes into the set $P_{>\ell}(\mathcal{S})$.

Let Δ_ℓ be the contribution to Δ due to flow along the paths in P_ℓ, and $\Delta_{>\ell}$ the contribution due to flow along the paths in $P_{>\ell}$. Then,

$$\begin{aligned}
\Delta &= \Delta_\ell + \Delta_{>\ell} \overset{(a)}{=} \Delta^\ell_{v_0} + I(v_0|P_{>\ell}),\\
&\overset{(b)}{\leq} \Delta^\ell_{v_0} + \alpha I(\mathcal{S}|P_{>\ell}(\mathcal{S})),\\
&\overset{(c)}{\leq} \Delta^\ell_{v_0} + I(\mathcal{S}|P_{>\ell}(\mathcal{S})),\\
&\overset{(d)}{\leq} \Delta^\ell_{v_0} + \sum_{s\in\mathcal{S}} \Delta^\ell_s \overset{(e)}{=} \delta(\ell) \overset{(f)}{\leq} \alpha^\ell p_i.
\end{aligned}$$

(a) follows from the definitions of $\Delta^\ell_{v_0}$ and influence; (b) follows from Lemma 3 and (c) because $\alpha \leq 1$. (d) follows because the paths in $P_{>\ell}(\mathcal{S})$ are all of length ℓ, so $P_{>\ell}(\mathcal{S})$ is a subset of all the paths of length ℓ that terminate in $\mathcal{S}$; (e) follows from the definition of $\delta(\ell)$, since $\mathcal{S} \cup v_0 = N_\ell(v_i)$; finally, (f) is an application of Lemma 2. Equality occurs *iff* $\mathcal{S}$ is empty, and all paths from v_i are of length ℓ, ending at v_0. Certainly, the optimal value of ℓ is 1, and so we have the following theorem[4].

Theorem 3 *Δp^i_0 is maximized if and only if the* only *edge from v_i is to v_0. This is independent of all the other edges in the graph, in particular independent of the edges from the other v_j.*

Theorem 3 directly implies the following result,

Corollary 1 *The direct individual attack is optimal.*

A related issue is whether the direct individual attack also maximizes the rank (as opposed to the pagerank) of the victim. This question is not immediately answered by Theorem 3 since the actual rank depends on the *relative* pagerank of v_0 with respect to the other nodes, and not the absolute pagerank of v_0. We will now show that the rank is also maximized by the direct individual attack.

Suppose that some other attack X maximizes the rank of v_0. This means that for some node u, $\bar{p}^I_{v_0} \leq \bar{p}^I_u$ and $\bar{p}^X_{v_0} > \bar{p}^X_u$ (I denotes the direct individual attack). We show that such a situation can never occur, leading to the following result.

Theorem 4 *The direct individual attack maximizes the rank of v_0.*

3.3 The Optimal Disguised Attack

We now consider the situation in which the attackers wish to maximize the magnitude of their attack on v_0, but they wish to disguise the attack by not pointing directly to the victim. In such an attack, the anchor text will be associated to the victim, hence we assume that the victim already has a high prominence with respect to the anchor text. The specific disguise constraint we consider is that for every attacker, the shortest path to the victim should have length at least $\ell \geq 1$.

Consider attacker v_i. In any attack, some amount of pagerank flows from v_i to v_0. In any directed graph, we define $f(u;v)$, the *forward value* of vertex u with respect to vertex v, to be the fraction of u's pagerank that flows to v along paths with v as

[4]An alternative proof of this theorem using the Markov chain approach can be given using a generalization of the result in [5], where it is shown that adding the edge (i,j) can only increase the rank of j.

terminal node but not as intermediate node. Thus, for example, $f(v;v) = 1$. Since the fraction of u's rank that makes it to v can be obtained by multiplying the fraction flowing to each neighbor with the fraction flowing from that neighbor to v, we obtain the *forward equation* for the forward values $f(u;v)$:

$$\begin{aligned} f(v;v) &= 1, \\ f(u;v) &= \frac{\alpha}{outdeg(u)} \sum_{(u,w)\in E} f(w;v). \end{aligned} \quad (3)$$

The forward equation (3) is similar to the pagerank equation (1) and can be solved by an iterative algorithm similar to the PageRank iteration [12].

For every vertex u (not an attacker), we consider the edge set $E_u = E \cup (v_i, u)$, which defines a new directed graph in which the edge set is augmented by a single link from the attacker to u. For this graph, we can compute the forward value $f_u(w; v_0)$ of any vertex w with respect to v_0. We define the value $V_i(u)$ of vertex u to attacker v_i by

$$V_i(u) = f_u(v_i; v_0).$$

By Lemma 4 the optimal attack is the one that maximizes the flow of pagerank to v_0, which means that v_i should point to the node u satisfying the "disguise constraints" that maximizes $V_i(u)$. Arguments similar to those that led to Theorem 3 give

Theorem 5 *The optimal disguised attack for a single attacker v_i is a* single *link to the vertex u, at distance $\ell - 1$ from v_0, which maximizes $V_i(u)$.*

Unfortunately, the maximizing node $V_i(u)$ need not be the same for different attackers – the disguise constraint introduces dependencies between attackers, i.e., the optimal attack for a particular attacker may depend on what the other attackers do. In particular, it is no longer the case that each attacker using its optimal disguised individual attack will maximize the magnitude of the disguised attack if the group of attackers act jointly. The following example with two attackers and $\ell = 2$ illustrates the issue.

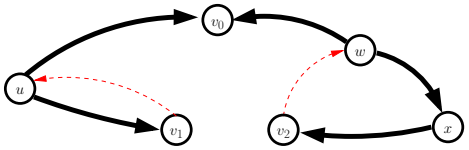

(a) Optimal individual attacks.

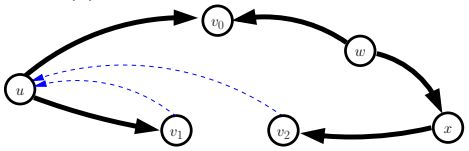

(b) Optimal joint attack.

The optimal attack for v_1 is to point to u for v_2 is to point to w (red dotted arrows in (a)). However, if both attackers attack, then they should both point to u. This is generally true,

Theorem 6 *There is an optimal joint attack in which all attackers point to the same intermediate node u which is distance $\ell - 1$ from v_0.*

A detailed comparison of the optimal joint attack with the greedy strategy in which the attackers each adopt their individually optimal attacks is beyond the scope of this current paper.

4 Experimental results

In this section, we give some preliminary experimental results that quantify the effectiveness of Google bombs in various environments. There are four main degrees of freedom we explore: the nature of the graph, including its connectivity or edge density; the prominence (pagerank) of the attackers; the prominence of the victim; and, the value of α.

We ran our experiments on three types of graphs: *Random* is an Erdös-Reyni type ($G(n,p)$) random graph with edge probability p; BA (Barabási-Albert) is a preferential-attachment random graph with 5 outgoing edges per vertex [2]; (Such graphs are known to have power-law in-degree distributions, and since we add the vertices sequentially, there are no cycles.) $MWDTA$ is a modified "Winner's don't take all" random graph in which every node has at

least one out-going edge method [13]. (Such graphs are known to model certain characteristics of the world wide web graph such as power-law in and out-degree distributions.). The main difference between $MWDTA$ and BA random graphs is that in $MWDTA$, a larger number of nodes will have significant indegree, whereas in BA a few nodes have very large in-degrees. In order to make fair comparisons, we normalize graphs from different random graph models (*Random*, BA or $MWDTA$) to have the same expected number of edges.

First, we generate a random graph with 1,000 nodes, and randomly select 10 attackers and a victim. We then remove outgoing edges from the attackers and perform a pagerank computation, obtaining:

p_0, the page rank of the victim;
p_A, the average pagerank of the attackers;
$f_p(p)$, the pagerank distribution in the graph;
σ_p, the std. dev. of the pagerank distribution.

We only show results for two of the attacks described in Section 3.1: the optimal direct individual attack I, and the cycle attack C (the results for other suboptimal attacks are similar). Each attack is repeated a number of times on randomly generated graphs to increase the statistical significance of the results. We use the following measures of success for attack X,

$G(X) = \text{Gain} = \Delta p_0^X / p_0$,
$\bar{G}(X) = \text{Normalized Gain} = \Delta p_0^X / \sigma_p$,
$D(X) = \text{Discrepancy Factor} = G(I)/G(X)$.
$\bar{D}(X) = \text{Normalized Discrepancy} = \bar{G}(I) - \bar{G}(X)$.

The pagerank distribution $f_p(p)$ generally affects the effectiveness of an attack. Figure 1(a) shows pagerank distributions for the various random graphs. As can be seen, *Random* has a (near) Normal distribution, compared with BA and $MWDTA$ which have power-law type distributions in which $MWDTA$ appears to have a slightly fatter tail than BA.

Some detailed results on the effectiveness of the attacks are shown in Figure 1: (b) shows how connectivity (number of edges) in *Random* graphs with different p affects the attack; (c) shows different graph types; (b,c) show the dependence on the prominence of the attackers, and (d) on the prominence of the victim; (e) shows the dependence on α; and, (f) shows some results for the *rank* (as opposed to the pagerank). We give a summary of the results below.

Higher Density: All attacks decrease in magnitude (new edges have little additional effect when the graph is already dense).

Graph type: Prominence of attackers has (by far) the largest impact in *Random* graphs, then BA and $MWDTA$. (Pageranks in *Random* graphs are "concentrated" around the mean, so any bias in the victim's pagerank results in it becoming extreme. This is less so for BA and even less so for $MWDTA$.).

Higher Prominence of Attackers: Stronger attack.

Higher Prominence of Victim: Attacks become less effective and $D(C)$ decreases (diminishing returns).

Lower α: $D(C)$ increases (it is more costly to divert from the individual attack).

Rank: For random graphs, an attack usually results in a top ranking for the victim, which is not usually the case for BA and $MWDTA$ graphs.

5 Discussion

We have shown that the best attack is the direct individual attack, in particular: *any* organized structure among the attackers reduces the impact of the attack; links that cycle back to attackers in an attempt to boost their pageranks are detrimental. The discrepancy between the optimal individual attack and suboptimal attacks can strongly depend on the graph type through the initial pagerank distribution. Our results indicate conditions that offer resistance to rank manipulation: dense, power-low type graphs in which victims already have high rank, attackers have low rank and α is small. Our analysis has been focused on increasing a page's rank (pagerank manipulation) in the entire graph, i.e., the victims rank is increased for *every* query. The underlying model is that the query identifies a set of nodes (based on text and anchor text), which defines an induced subgraph of the original graph. However, the nodes are ranked according to pagerank in the original graph. This model has the feature that pageranks do not need

to be recomputed for the specific query. An alternative approach is to order the nodes with respect to the pageranks in the induced subgraph (hence these pageranks would need to be recomputed for every query). Such a model would mean that one attempts to boost the pagerank with respect to a specific query and not others. Our analysis does not apply to this model, and it is no longer true that the optimal attack is the direct individual attack. The following example (with a single attacker) illustrates the issue.

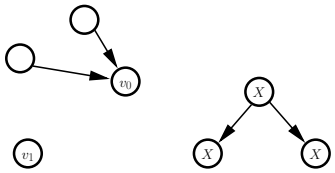

(a) Original graph.

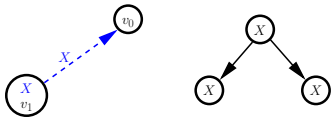

(b) Direct attack.

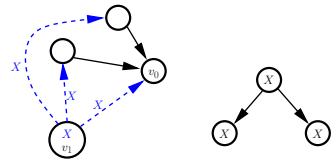

(c) Indirect attack.

In (a) we show the original graph, where X will be the query text and the attacker wants to boost the rank of v_0 with respect to X. In (b) we show the subgraph induced by the direct attack, where the attacker places X in its page as well as in the anchor text of the link. In the resulting induced subgraph, the rank of v_0 is not the highest. The benefit of the non-direct attack in (c) is that other nodes that point to v_0 get included into the induced subgraph. Thus while the flow of rank from v_1 to v_0 is decreased, this is more than compensated for by the additional rank contribution from the newly included nodes. A better attack would arise if v_1 added another link to v_0. In fact for any attack in which v_1 has k links to v_1, a strictly better attack with $k+1$ links is possible. In this example, there is no optimal attack. In general, we can formulate this notion by saying that the attacker should add the minimum number of links to all nodes with paths to the victim which do not contain the query text, and hence would not be included in the subgraph. The attackers should then place as many parallel direct links as feasible. The end effect is to include all nodes with paths to the victim with a minimum diversion of page rank. Of course, such a huge attack is not very practical, and an interesting question is to consider the optimal attack under this model when each attacker has a fixed budget of links.

The PageRank algorithm favors attacks from groups that are not well connected, which makes it harder to detect the attack, and accountability in such an attack formation becomes an issue: who is responsible for the attack? Different variations of the PageRank algorithm may suffer a similar fate if they propagate the pagerank in a similar way (for example Topic-Sensitive PageRank [8], provided that the attacking group is considered relevant to the query). In order to avoid such a fate (a dilemma faced by any ranking method open to manipulation by small groups), either one must change the ranking function or somehow exclude the attacking group from the search engine's database. While such an approach is a reasonable way to deal with private companies attempting to manipulate rankings based on their own views, it is not very democracy-friendly to arbitrarily remove certain pages from a search engine.

As discussed in [6], the PageRank algorithm makes certain assumptions about the user navigation patterns and the web structure that may not apply to the Web anymore. [6] considers the effect of dangling nodes in the pagerank computation and provides methods to adjust for them. They also point out that users will rarely (if ever) navigate to one of several billion pages uniformly – they may not even know that these pages exist. In fact, users generally start from known sites and navigate from there. Hence, random navigation is more likely to bring them to one of these "anchor" sites. The HostRank algorithm [6] uses this assumption to choose a set of

anchor sites, and they show that such an approach is more resistant to attacks. A related issue is that of navigation along links from a site. One is more likely to trust a link on a highly ranked page, and one is more likely to follow a link to a highly ranked page. For example, it might be *much* more probable to follow one of the links from a search engine or a news Web site than a regular web page. The probability to navigate from a page in the PageRank algorithm is independent of a page's rank, and the link one selects to navigate is random. A plausible alternative is that the probability to navigate from a page should be proportional to the page's pagerank, and the probability to use a particular outgoing link is proportional to the pagerank of the destination page. Such a navigation model would lead to an equation (analogous to (1)) of the form

$$p_i = \kappa\alpha p_i \sum_{(v_j,v_i)\in E} \frac{p_j^2}{\sum_{(v_j,v_k)\in E} p_k} + \frac{1-\alpha}{N}.$$

More effort could be spent how the transition probabilities generally affect the pageranks and their manipulability. [6] discusses such issues for nodes with unknown outgoing links and [15] uses the amount of traffic flow through the nodes to model the transition probabilities. It would be interesting to see what the optimal attack with such ranking algorithms is. In short, objective methods for the selection of the anchor sites or more plausible navigation models deserves closer examination. One must also bear in mind (see for example [6]) that the computational complexity of the algorithm is also an important practical consideration for any ranking algorithm.

Other factors, which we do not study here, might be significant to the success of an attack. [7] argues that anchor text pointing to a page gives information regarding the subject matter of that page, and relationships between different pages. For example, Google may consider both the pagerank and the frequency of keywords in links pointing to a page when computing the score of the page. Google bombs in the past used the same keywords when pointing to the attacked page, i.e., the bombing links were correlated in that they all had the same keywords, whereas in general, links pointing to a website would not display such a correlation. If some linear combination of these two factors is then used in the final score, it will favor attacks over the natural Web behavior. If some small group of sites use a specific keyword to point to a victim, it is unlikely that this groups's sites are unrelated, and one could (for example) add pseudo-links among these sites, since the expection would be that they participate in some group structure. As our results show, these pseudo-links will reduce the magnitude of the attack. One could go so far as to say that if after the addition of such pseudo-links in the graph, the pagerank distribution does not change significantly, then the ranking algorithm should be more resistant to manipulation.

The analysis of the optimal attack structure provides a new tool for looking at resistance to link manipulation. Such metrics and an understanding of optimal attack formations for other algorithms should be fruitful directions for future work.

Acknowledgement. We are grateful to Mark Goldberg for his initial feedback on this work.

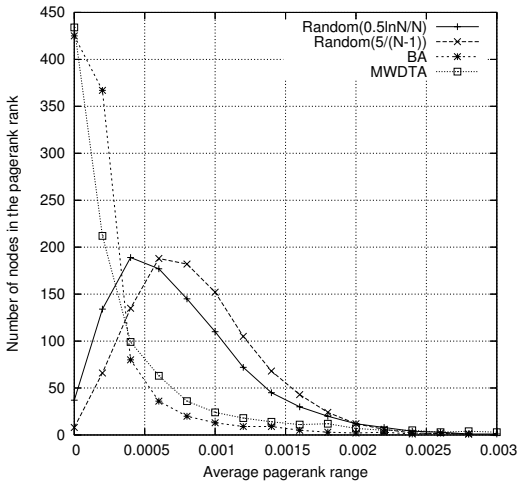

(a) Pagerank distributions of different graphs

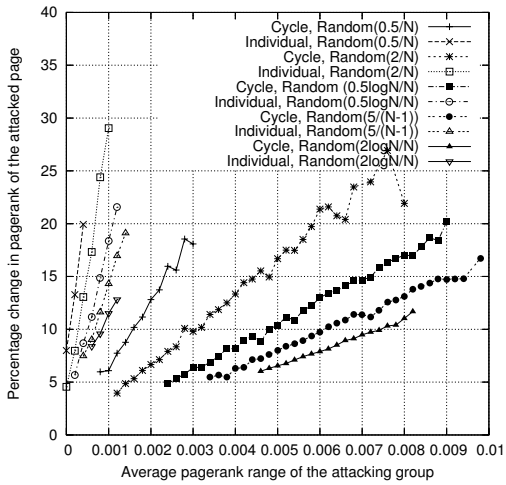

(b) Graphs with different edge densities.

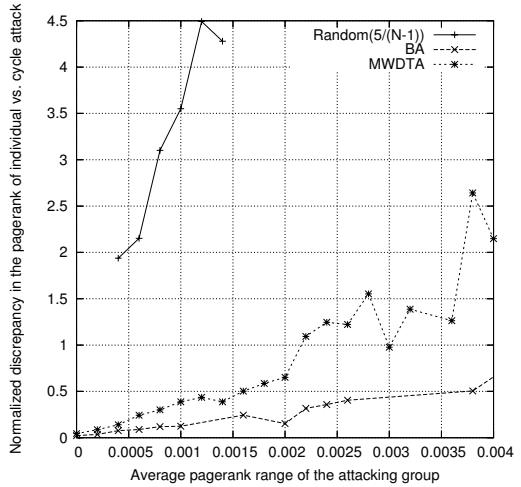

(c) Attacks in different graph types

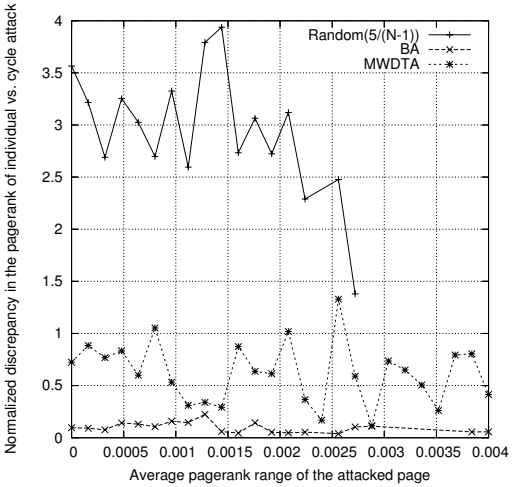

(d) Dependence on pagerank of the victim

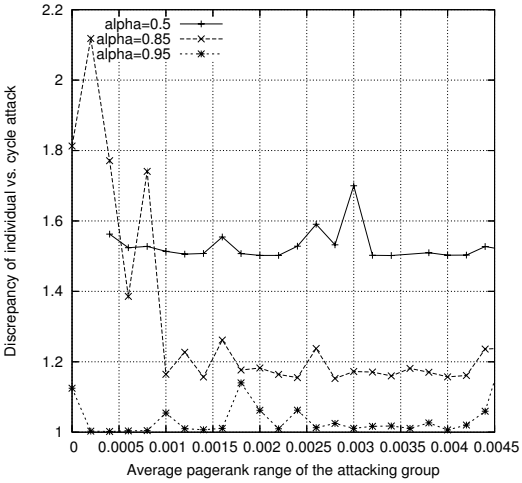

(e) Dependence on alpha (α) for MWDTA graphs

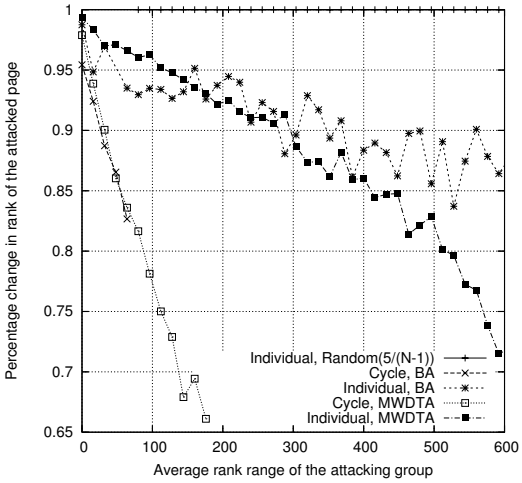

(f) Change in rank of victim

Figure 1: Experimental Results for $n = 1000$.

References

[1] S. Adali, T. Liu and M. Magdon-Ismail, *An analysis of optimal link bombs*, CS Department Technical Report, TR#05–11, RPI, Troy, NY, 2005.

[2] A.-L. Barabàsi and R. Albert. *The Emergence of Scaling in Random Networks*, Science, 286, 1999

[3] S. Brin and L. Page. *The Anatomy of a Large-Scale Hypertextual Web Search Engine*, Proc. 7th WWW Conf., 1998

[4] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. *Graph structure in the web*, Computer Networks, 33(1-6): pp. 309–320, 2000

[5] S. Chien, C. Dwork, R. Kumar, D. Simon, and D. Sivakumar, *Towards exploiting link evolution*, Workshop on Algorithms for the Web, 2002.

[6] N. Eiron, K. S. McCurley, J. A. Tomlin, *Ranking the Web Frontier*, Proc. 13th WWW Conf., 2004.

[7] N. Eiron and K. S. McCurley. *Analysis of anchor text for web search*, Proc. 26th SIGIR, 2003.

[8] T. H. Haveliwala. *Topic-sensitive pagerank*, Proc. 11th WWW Conf., 2002.

[9] http://www.microcontentnews.com/articles/googlebombs.htm

[10] T. McNichol. *Engineering Google Results to Make a Point*, NY Times, Jan 22, 2004.

[11] A. Y. Ng, A. X. Zheng, and M. I. Jordan, *Stable algorithms for link analysis*, Proc. 24th SIGIR 2001.

[12] L. Page, S. Brin, R. Motwani, and T. Winograd. *The pagerank citation ranking: Bringing order to the web*, Stanford University Database Group TR, 1998.

[13] D. M. Pennock, G. W. Flake, S. Lawrence, E. J. Glover, and C. L. Giles. *Winner's Don't Take All*, Proc. National Academy of Sciences, 99(8), 2002.

[14] D. Sullivan. *Google's (and Inktomi's) Miserable Failure*, searchenginewatch.com, Jan 6, 2004.

[15] J. Tomlin, *A New Paradigm for Ranking Pages on the World Wide Web*, Proc. 12th WWW Conf., 2003.

[16] H. Zhang, A. Goel, R. Govindan, K. Mason, and B. Van Roy, *Making eigenvector-based reputation systems robust to collusion*, Workshop on Algorithms and Models for the Web Graph (WAW), 2004.