Media Monitoring Using Social Networks

Tony White
Carleton University
1125 Colonel By Drive
Ottawa ON K1S 5B6

arpwhite@scs.carleton.ca

Wayne Chu
Carleton University
1125 Colonel By Drive
Ottawa ON K1S 5B6

Amirali Salehi-Abari
Carleton University
1125 Colonel By Drive
Ottawa ON K1S 5B6

asbari@scs.carleton.ca

ABSTRACT

With the rapid rise in the number of weblogs, or blogs, on the World Wide Web (WWW), there is a growing need to be able to quickly search for discussion on specific topics. While keyword searches using tools such as Google [4] or Technorati [18] can yield useful results, we run into the problem of having to enter contextualizing keywords to filter out unwanted and irrelevant search results. This has the unfortunate consequence of making the search process more complicated and possibly filtering out search hits that we would typically want. This paper outlines an approach to narrow search results to only relevant hits, while allowing for general keyword queries. Since the blogosphere constitutes a social network, the solution, BlogCrawler, attempts to use the properties of social networks to narrow the focus of search queries to only those blogs that the user is interested in. This paper presents an algorithm and empirical evaluation that exploits the social network implicit in blogs found on the WWW for the purpose of improving search on the Web.

1. INTRODUCTION

Increasingly today, corporations, governments, and private citizens are demanding to know what others are thinking. Corporations would like to know what the latest trends are in society today, governments need to know how policies are being received amongst the general population, and private citizens often are simply interested in knowing what their peers are thinking about. Weblogs, or blogs for short, are a recent phenomenon in cyberspace that has emerged which offers an incredibly useful collection of information for media analysis. Blogs offer individuals the ability to read what people on the Internet are thinking right now. Indeed, the nature of blogs, online journals written by individuals around the world that are updated frequently, often daily, present an unfiltered view of world, discussing whatever the author feels like talking about. The content found on blogs varies greatly from topics such as technology tips and tricks, politics, arts and entertainment, and even personal accounts of the author's daily life. The unedited, uncontrolled nature of this media means that there are no limits to what blogs can talk about.

The rapid growth in the number of blogs on the Internet also means that this phenomenon is not something restricted to a small subset of the most technologically literate individuals in society. This makes blogs highly relevant and an attractive target for data mining. While, only a handful of blogs initially existed in 1998, the numbers have grown exponentially over the years to the point that they now number in the millions, according to leading blog tracker, Technorati [12], [18]. Many blogs [14] also have high readership numbers, such as the popular technology blog, Slashdot, which has an audience numbering in the hundreds of thousands [13]. With millions of blogs and their associated authors, or "bloggers", expressing ideas, the amount of information one could find in this subset of the World Wide Web is clearly immense.

More important to this paper, however, is the fact that blogs are not simply self-contained journals on the Web. Like any website, blogs contain links to other sites – this includes other blogs. Bloggers will link to others with similar interests. These linkages allow sites to interact with each other, forming an online community. What emerges from this is the formation of social networks within the so-called blogosphere [9], creating order amongst the chaos of the web. It is this observation that motivates the algorithms and their evaluation described in this paper. As Section 3.0 of this paper outlines, the power of social networks is great, so exploiting the inherent social networks in the blogosphere is something of great interest.

2. PROBLEM DESCRIPTION

The project described in this paper empirically evaluates if we can search the blogosphere to see what people are saying about any particular topic. We will attempt to harness the power of the blogosphere determining if useful, relevant search results can be returned using simple keyword queries. This may seem simple like a simple task since keyword searching is something that even the simplest search engines are capable of. However, with the ability for anyone, anywhere to start publishing a blog, finding the information one wants while minimizing the complexity of a query may not be as simple as searching every single blog for a specific search term. The meaning of a keyword is dependent not only on the dictionary meaning of the word, but also the meaning the human user placed on it and the context in which it is found. For example, if one were to want information regarding the recent Canadian federal budget, a simple search for the term "budget" would be inadequate. Searching for the term on the popular search engine Google produces over 93 million results – far too many for anyone to sift through. Even restricting the search to blogs using the blog search engine Technorati produces over 180,000 results with posts about a disparate range of topics. The technique for conventional searches, then, is to expand the keyword search to provide the needed context. For example, searching for the terms "Canadian federal budget" on Google reduces the number of hits to 3.2 million. Table 1 outlines the results of various keyword searches performed on the blog search engine Technorati. As we can see from the table, providing context clearly narrows the amount of hits returned. The trade off, however, is that search queries must become more complex.

Table 1: Search Results of Various Keyword Queries

Query	Hits	Query	Hits
budget	193,061	canadian federal budget	2211
senate	163,333	canadian senate appointments	152
hotel	440,388	moby cd hotel	326

Source: http://www.technorat.com on 2005/03/31at 7:04pm.

The question that this paper attempts to answer, then, is whether or not it is possible to return concise, relevant search results from the blogosphere while minimizing the complexity required in the actual search query. To do this, we will use the power of social networks to limit the scope of search queries such that a simple keyword such as "budget" will return only pages that the user will want to see. More specifically, if we search only within blogs located within -- to continue our example -- a social network of Canadian political sites, then the need to provide contextualizing keywords in the search query will no longer exist. This is analogous to searching for books only within a specific genre. It may be helpful first, however, to outline what a social network is in the context of computer networks and the web.

3. SOCIAL NETWORKS AND BLOGS

In the physical world, a social network is a "network of friendships or other acquaintances between individuals" [3]. As friends or acquaintances, these individuals often share common interests and backgrounds with each other. At a basic level, everyone participates in a social network, defined by the everyday interactions one goes through. The concept is one that spans disciplines such as sociology, political science, and in our case, computer science, since in all cases the nature of how we form and keep human relationships is of great interest. We can easily transfer the idea of the social network to that of computer networks, and more specifically, to the blogosphere. In this case, each blog acts as an individual, or a node in the network, and the hyperlinks on each blog act as the connections between the nodes. Figure 1 demonstrates the basic structure of a social network in the blogosphere as an undirected graph, showing how the network of blogs is analogous to a physical network of friends and classmates. We can further transfer the concept to computer networks when we observe that the people who use computer networks "have social relationships with each other that are embedded in social networks" [19]. From this, we can infer that the network of blogs within the blogosphere has an inherently intelligent human-based organizational structure. This structure has several important properties that we will apply here to solve our search problem.

The first significant property is that of "community structure" in which we see similar nodes densely clustered together within a wider network. In the context of a physical social network, these clusters of nodes represent social groupings of those with common interests or backgrounds [3]. Applying this to the blogosphere, and treating the hyperlinks between blogs as connections in the network, we would expect blogs to be clustered together such that they all cover similar types of topics. Therefore, we can infer that if we were to start at a specific blog and search within it and it's neighbours, we would be predominantly searching blogs of a similar nature and type, providing the context that we were seeking earlier in our problem description.

The second property is the "small world effect", which states that the distance between two vertices in any network is short. Indeed work has suggested that this property is "pervasive in networks arising in nature and technology, and a fundamental ingredient in the structural evolution of the World Wide Web" [10]. The basis of this is an experiment performed by Stanley Milgram in which complete strangers in Nebraska were tasked with getting a letter to a stockbroker in Boston. These strangers could only pass the letter on to someone with whom he or she was on a first-name basis. Milgram found that the average number handoffs required for a letter to be received by the stockbroker

were only about six [15]. If one were to take each handoff of the letter as a connection in a network, we can see that within a social network of any sort, one can expect that the distance between any two nodes would be small. As a corollary, this means that a web crawler would only need to crawl very few levels before a sizable amount of the social network is covered. Moreover, Girvan and Newman discuss that many networks display a property of transitivity, in that nodes which share a common neighbour are likely neighbours of one another [3]. As will be discussed later in Section 3.1 this makes extracting a community of blogs from within the greater blogosphere much less complex since the depth one must crawl to retrieve a suitable number of blogs is low.

A third property is the idea of trust relationships in the social network. When a human decides to form a relationship with another, this forms a social exchange. In one regard, the exchange occurs to fulfill a purpose of some kind. In another regard, "exchanges involve investments, gains and losses of time, money, energy, emotions, expectation, and many other energetic and motivational elements" [16]. Put simply, when someone in a social network creates a connection with another individual, then there is an implicit recognition that making that connection was worthwhile. In other words, an individual trusts the other enough to make a connection. Within the blogosphere, connections are made in the form of hyperlinks to other blogs. The fact that a blog has linked to another means that there is some value in the other blog. Extending this idea, if a blog is linked to by many people, then, theoretically, this means that the blog is found to be worthwhile by many, improving its level of trustworthiness. Kleinberg [21] recognized this with identification of hubs and authorities. This concept, as will see later in the paper, is important in determining which blogs are more valuable than others.
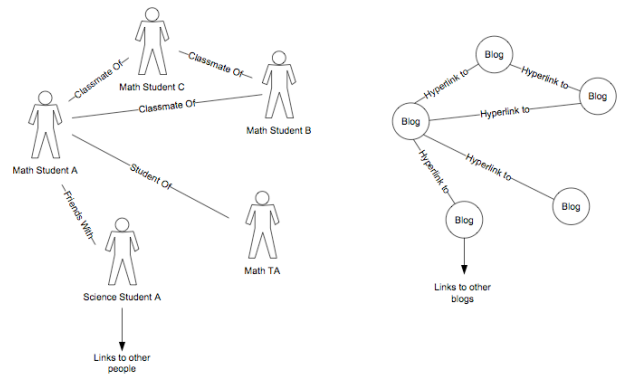

Figure 1: Graph of a Social Network

4. SOLUTION
4.1 Overview

Let us return, now, to the original problem outlined in this paper. Namely, that of narrowing search results while maximizing the amount of generalization possible within a keyword search query. To accomplish this, we will implement a blog search engine, called BlogCrawler which will use the inherent social networking properties of the blogosphere, ensuring that to the end user searching for information is as simple as using any other search engine available on the web. In the perspective of the end-user, he or she will still be required to enter at least one keyword to search on. However, we will eliminate the need for context specifying keywords, such as "Canadian" and "federal", to use our previous example, by limiting the number of sites that we will

search to those of a specific topic. Instead of searching millions of sites for a specific keyword, we will only search hundreds.

This obviously leads to the question of how we determine which sites to include in our query and which to ignore. This is where we use the properties of social networks to aid us. We know that because of the property of community structure, neighbours of a blog will be of similar type. We also know that due to the small-world effect, the members of the social network of blogs we want to search will all be situated near each other in terms of number of links traversed to get from one to another. Finally, in the context of our problem, we know what types of blogs we are interested in. For example, we will know beforehand that we want information about Canadian politics. Therefore, in our solution we will pre-select a small number of "expert" blogs deemed to be representative of the type of blogs we want to search. We will treat these selected blogs as the root nodes of the social network being traversed, and then crawl through the blogs connected to them. The end result is that after traversing only a few levels of links, we will have indexed a sizable number of blogs, most of which should be similar in type to the root blogs.

Once we have a collection of blogs to search from, then finding the information we want is simply a matter of entering a general keyword and performing a text-based search of the content of the blogs collected. If we were successful in limiting the type of blogs in our collection, then the results returned should be limited to only those topics that we are interested in. There is one caveat. Because of the uncontrolled nature of the blogosphere, anyone can write anything they want. As well, some blogs have more relevant information and are more trustworthy than others. This is where the idea of trust relationships within a social network comes into play, higher trust indicated by the authority of the site [21]. Each blog will be assigned a rank, similar to that of the PageRank score given to websites on Google [13]. This rank will represent a blog's level of trustworthiness. Hence, the results of a keyword search will be sorted such that those sites with the highest rank are situated at the top of the search results, ensuring the most relevant and trustworthy hits are the first ones the user sees.

There are several components implemented for such a solution. Those are a web crawler, a ranking module to calculate the expert level of blogs, and a front-end user interface for the user to access when performing search queries. Figure 2 outlines the architecture of the system.


Figure 2: System Architecture

4.2 Software

BlogCrawler is implemented in Java, including a Java Server Pages (JSP) application for the front-end web-based search engine. Java was chosen for a variety of reasons, including the fact that it is cross- platform compatible, allowing all types of machines to use the same code base. This simplifies the task of ensuring that the code can run on the widest range of machines as possible. In addition, the standard Java library also includes tools to easily connect to, retrieve, and tokenize HTML web pages, which is helpful in the implementation of the search engine.

Another deciding factor to use Java was also the decision to use the Apache Lucene library, a full-featured text search engine. Lucene allows developers to efficiently store and search large amounts of any type of data, including web sites. Since a large portion of the project involves performing keyword searches on text, it was felt that using Lucene would ensure the overall efficiency of the system. Because the engine is written entirely in Java, and the fact that every element of the project would need to access the index, it was felt that the use of Java was the most appropriate choice of environment.

4.3 Web Crawler

The web crawler is responsible for gathering the blogs which we want to search and is the most complex component in the system. Unlike standard web crawlers, the requirements of the system mean that it cannot blindly follow every single link on a blog. Since the blogosphere is only a subset of the greater World Wide Web and individual blogs often post links to sites that are not blogs, doing so would lead to a massive amount of web sites crawled that are outside the scope of our search parameters. Clearly, the web crawler needs to be intelligent enough to crawl only those sites we want. To solve this issue, the crawler needs to be equipped with a validation module that verifies whether or not a web page is a blog. Only if the validation determines that a page is a blog will it allow the crawler to index the site. The end result should be that only those sites situated within the blogosphere will be searched.

The basic algorithm for web crawling is a simple one. Simply access a site, store its contents in the index, extract the links on the site, and recursively crawl those links. Note, however, that due to the nature of the web, actually using a recursive algorithm would create an exponential number of instances of the crawler. We can easily solve this problem by using a queue instead as seen in Algorithm 1 [1].

```
Algorithm Crawl(firstSite)
    Queue queue := new Queue;
    queue.enqueue(firstSite)
    While queue not empty
        currSite := queue.dequeue();
        siteContents := getContents(firstSite);
        outgoingLinks := getLinks(firstSite);
        index(currSite, siteContents);
        queue.enqueue(outgoingLinks);
    Loop
```

Figure 3: Algorithm 1

As we noted earlier, we cannot follow every single link. There are several issues when programming a web crawler that we need to be cognisant of. First, we need to ensure that the web crawler does not continue to run for an indeterminate amount of time. This is a distinct possibility with millions of blogs on the

web online today and the number of web hosts doubling every year. To resolve this issue, we simply need to limit the depth that our web crawler will crawl. For example, assume that the initial list of blogs we crawl represent depth zero. Then every blog linked from those at depth zero would be depth one, those linked from depth one will be depth two, and so on. Therefore, we simply place a condition on our crawler to stop indexing blogs that exceed a user-defined depth.

The second issue we need to address is that of ensuring that only blogs are crawled. We must remember that the blogosphere is not self-contained within the World Wide Web since blogs not only link to other blogs, but to other websites that we do not want to index. If we allow the crawler to exit the blogosphere into the greater web, then the chances are slim that the crawler will return to where we want it to. This problem can be addressed by validating each page the crawler retrieves before indexing it. If the page validates as a blog, then the crawler will index it, extract the links on the page, and continue crawling. If it does not validate it, the crawler will ignore the page. Section 4.4 goes into further details on how the validation works.

Finally, we need to ensure that we do not index pages multiple times, since most blogs, and web pages for that matter, maintain a many-to-many link relationship with other blogs. This is resolved by maintaining a list of websites already crawled and skipping those links which lead to previously traversed pages.

From our basic algorithm, then, we now have a more intelligent and efficient crawling algorithm, as seen in Algorithm 2.

```
Algorithm IntelligentCrawl(rootSites, maxDepth)
    List visited = new List;
    Queue queue = new Queue;
    queue.enqueue(rootSite)
    while queue not empty
        currSite := queue.dequeue();
        if currSite.depth <= maxDepth then
            visited.add(currSite)

            if not visited.contains(currSite) then
                if isBlog(currSite) then
                    siteContents := getContents(currSite);
                    outgoingLinks := getLinks(currSite);
                    index(currSite, siteContents);
                    queue.enqueue(outgoingLinks);
                end if
            end if
        end if
    loop
```

Figure 4: Algorithm 2

Looking at each individual blog to be indexed, we now have to determine what properties of an individual blog we need to store. We have already mentioned that we will be using the Lucene library to index the pages that we crawl. As part of its implementation, Lucene allows us to index any number of fields with whatever content we wish. The obvious fields to index are the ones the end-user is interested in, namely the address, title, and contents of the page. These are not the only things that need to be indexed, since we also need to store attributes that relates the page to its position within the social network. Our discussion of social networks identified two key aspects of blogs that we are interested in. Firstly, we need to know what other blogs the page is connected to. To that end, we also index the complete list of outgoing links contained on the page. Secondly, we need to know the rank of the page to determine which blogs are the most trustworthy. Therefore, we will also index the page's rank. Since a page's rank is dependent on the rank of other pages in the index,

we will not be able to calculate the rank during the web crawl, so we index the page with an initial rank of zero. All of this is accomplished with the PageIndexer.

Now that we have the overall structure of the crawler module designed, we turn to ensuring that the module is as efficient as possible. Like any algorithm that we design, we are particularly concerned with time and speed efficiency and memory usage. In both cases, implementing the web crawler created many challenges in ensuring that our algorithm was as efficient as possible.

With respect to speed efficiency, we needed to ensure that crawler indexed as many pages as possible in a minimal amount of time. Particularly troublesome was the fact that as the crawler traversed deeper into the network, the number of links queued to crawl grew exponentially in the same way that the number of leaves on a tree grows exponentially at each height. A high performance data base was used to store the results for each page crawled.

4.4 Blog Validator

To validate whether a site is a blog or not, we created a BlogValidator which is designed to filter out unwanted sites. We accomplish this by recognizing the fact that blogs share common formatting characteristics, including a consistent sequence of date-entry pairs [13]. Particularly, we note that the date on each entry is consistently formatted in terms of date expression (such as "dd/mm/yy") and formatting style (such as font size, bolded). The validator takes advantage of this by analysing the structure of an HTML document, extracting the dates on the page and determining if the sequence of dates consistent with that of a blog. Specifically, the validator determines that a site is a blog if and only if:

1. There exists a sequence of dates, spaced out by a user defined minimum number of characters.

2. The sequence of dates is ordered in ascending or descending order.

3. The HTML tag sequence surrounding each date instance is uniform for all entries in the sequence.

To accomplish this, the BlogValidator extracts all the dates on a page that match one of a number of predefined regular expressions. If a date follows another date within a specified number of characters, however, it is ignored. Once extracted, the dates are sorted into bins based on the regular expression that the date matches. For example, the dates "2005/02/14" and "2004/12/01" would go in the same bin, having matched against the date format "yyyy/mm/dd". The bins are then further split by the HTML tag sequence the date was found in. For example, a date may be found following a tag sequence of "<div><span><p>". We then take the largest list of dates, and assume that the list represents the sequence of date entries for the articles on the potential blog. If those entries are found to be in ascending or descending order, then the page is determined to be a blog and the validator returns true.

Using this algorithm, we are able to detect blogs produced by any type of blog software and presented using any type of template. While thought was given to simply using a precompiled list of "acceptable" blogs, this approach of determining whether a page is a blog in real-time allows for much more flexibility in discovering lesser-known blogs. In fact, as will be discussed in

Section 4.1, this algorithm is remarkably successful in identifying blogs.

4.5 Page Ranker

After we use our web crawler to extract and index the blogs were are interested in, we now turn to ranking the blogs to ensure that the most important ones are given the most weight when performing keyword searches. Doing this prevents blogs that are relatively new and those that are unpopular from being returned ahead of the blogs that the end-user actually wants to see. We accomplish this by assigning a rank to each blog which represents its relative importance or trust level compared to other sites in the index. When performing a keyword search, blogs with higher rankings will appear first in the results listing, followed by those with lower ranks. This is not a new concept, as demonstrated by Google, which returns search results based on the rank of a page calculated by the highly successful PageRank algorithm [2]. So, rather than create our own algorithm, we will use a modified version of the PageRank algorithm to calculate the ranks for the blogs in our index.

PageRank is premised on the idea that the more important a web page is, the more other pages will link to it. Therefore, the more links to a page, the higher its rank is and the higher its importance. Mathematically, PageRank begins by assigning each page an initial rank of $1/N$, where N is the number of pages in the index. Let Nu be the out degree, or number of outgoing links, on page u, and let $Rank(p)$ be the rank of a page p. Also, let Bv be the set of all pages with a hyperlink to page v. The rank of a page, v, at iteration i is calculated as follows:

$$Rank_{i+1}(v) = \sum_{u \in B_v} Rank_i(u) / N_u \qquad (1)$$

Since the rank of a page is dependent on the rank of others, we iterate through all of the pages in the set of pages until the ranks stabilize to within a specified threshold. The rank vector that is calculated from this formula is calculated once and the results are used for every search query.

The PageRank algorithm, however, is susceptible to the problem of assigning pages with little actual authority a high rank simply because the page was heavily linked to (see [5]). This is a problem if we want to generalize keyword searches as much as possible while maximizing the relevancy of the search results. To overcome this obstacle, we bias the rankings by creating a "topic-sensitive" PageRank. As Haveliwala argues in [5], biasing the page ranking towards specific pages allows for personalization in the ranking. In this case of BlogCrawler, we will personalize the rankings so that blogs which we deem are experts or important will bias the ranking of all the blogs in the index towards them. Take the example of someone searching political blogs. If a user is specifically interested in conservative blogs, then it would make the most sense to bias the rankings towards those sites which present a conservative viewpoint. If the user is interested in liberal blogs, then using same index, the user can bias the rankings towards those blogs with a liberal viewpoint. This allows us to further narrow the search results to what the end-user wants.

Biasing the ranking algorithm is actually quite simple. Essentially, we want to ensure that those blogs which the user finds important have high scores. To do this, we modify the initial rank given to a page, such that a page has a rank of 1 if it is in the list of "expert" blogs, and a rank of 0 if the blog is not in the list. On each iteration, the PageRank is calculated as outlined in (1). The difference is that on each iteration, the initial rank given to the expert blogs is further diffused across the entire network, meaning that a blog's ranking is almost entirely dependent on its proximity to an expert blog.

Implementing this algorithm using the Lucene index is simply a matter of updating the appropriate rank field for each blog in the index. One final issue that needs to be resolved is that often times, there will be multiple pages from the same blog site indexed. To ensure that these internal linkages do not affect the final rank of a blog, we exclude outgoing links that point to pages on the same web host when calculating the rank. Once the ranks are calculated the index is now ready for keyword searching via a web-based search engine.

4.6 User Interface

To the end-user, searching through the blogs we have crawled should be as simple as using any other search engine. Modelled after a standard search box Figure 5 shows what the BlogCrawler search box looks like. The common element necessary present is a text box to enter the keywords the user wishes to search for. To search for information, all the user is required to do is enter his or her search query in the textbox, click Search and, as seen in Figure 6, wait for the search engine to display a list of appropriate blogs. Also present are two options to further refine the query. The user is able to enter a variety of keywords, separated by a space. Selecting "All Terms" will search through the index for items containing all of the keywords entered in the text box. Selecting "Any Terms" will return items that contain one or more of the keywords entered. In any case, if results are found, the search engine will return a sorted list of blogs, complete with an excerpt from the blog that matches the keywords entered.


Figure 5: User Interface

5. RESULTS

With implementation of BlogCrawler complete, we now move to evaluating the system. Recall that in the original problem statement, we wanted to know if it were possible to use the inherent social networking properties of blogs so that relevant hits would be returned using a generalized search query. Based on these parameters, we will focus on two main issues. We will first look at the web crawler's effectiveness in extracting the blogs we want, while filtering out everything else. Success in this regard is paramount since the success of the system requires that we actually have blogs to search from. Secondly, we will look at the actual results of keyword search queries performed on the index of blogs we constructed. We will determine whether or not we can retrieve relevant results from generalized queries and, if this is the case, relate this back to the idea of social networks.

5.1 Crawler Effectiveness

The web crawling portion of the BlogCrawler system can be evaluated in terms of two criteria. Firstly, how efficient was the actual crawling process. For this criterion, we can look at how fast web sites are crawled, and examine where the bottlenecks were during the crawling process. The second criterion is the accuracy

of the validator and whether or not we were successful in extracting only blogs.

In terms of the efficiency of the crawler, we were successful in producing a web crawler that maintained a consistent rate of operation in that regardless of how long the system was running, a web page was examined at least once every 10 seconds. Our crawler was able to index 1480 Canadian political blogs over the course of 2 days, not including sites that were rejected by the blog validator. This equates to approximately 30 blogs indexed per hour. While this performance rate is not poor, we were hoping to increase the rate of crawl. Section 5 outlines potential improvements that we can make to increase the web crawler's efficiency.
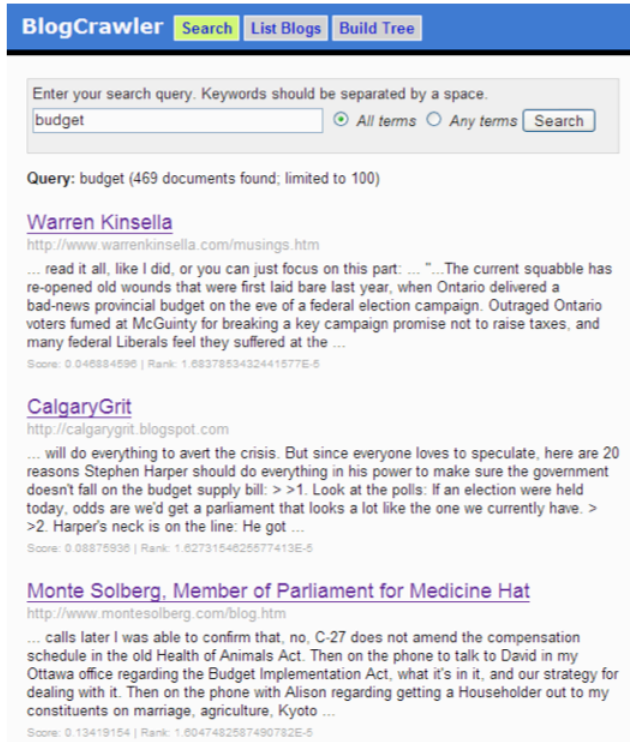

Figure 6: Search Results

Accuracy is the second criterion in which we can evaluate the web crawler. Indeed, for all other aspects of the system to work, we need to be confident that the sites being indexed are actually blogs. In this regard, we were very successful in implementing a validation module that more often than not, correctly identified web pages as blogs. These blogs include those based on standardized templates (e.g., http://calgarygrit.blogspot.com), and custom designed templates (e.g., http://www.freethought.ca). Using our list of blogs crawled, a random sample of 50 blogs revealed the validator incorrectly identified 4 web pages as blogs. In one case, there would be no easy way beyond natural language analysis to determine that the site was not a blog. That gives us a success rate of 94%, which for our purposes is more than adequate. Further random sampling generated similar results.

For the most part, then, our web crawler was successful in giving us a good base of blogs to analyze and search upon. Specifically, we were able to crawl a significant portion of the blogosphere that we were interested in.

5.2 Relevance of Search Results

Given that we have an acceptable base of blogs to search from, we now turn to evaluating the actual search results returned from BlogCrawler. In Section 1.3, we discussed how generalized search queries were inadequate in narrowing the scope of the returned results. Consequently, when we evaluate the effectiveness of our system, we need to look at how relevant the returned results are when using very general keywords. The main example used in this paper has been Canadian political blogs, so we continue to use this example in our evaluation of BlogCrawler. In our sample index, we begin by feeding two Canadian political sites (http://www.freethought.ca and http://calgarygrit.blogspot.com) into the web crawler and then allow it to run for two days. Afterwards, we rank the blogs using the two previously mentioned blogs to bias the rankings, then run the Updater module to ensure that all the pages in the index are recent. Finally, we perform various keyword searches using the BlogCrawler web application. For the purpose of comparison, we will also examine the results returned by the blog search engine Technorati (http://www.technorati.com) using the same keyword query.

Let us first examine a search on the keyword "budget", which happens to be a very general keyword that can apply to many situations. To satisfy the requirements of the system, we would like to only receive blogs that discuss issues surrounding the federal budget, an item prominent in the public consciousness at the time writing. More specifically, on March 31, 2005, when the historical search was run, the main issue was the federal government's decision to include environmental protection measures in the bill approving the budget. We therefore, would expect that the returned blogs would be discussing this issue. After running a search on BlogCrawler for the keyword "budget", we see that the returned blogs do discuss this issue for the most part. In contrast, a keyword search for "budget" on Technorati returns a plethora of blogs, with no consistent topic of conversation. Figures 7 and 8 show the results of the two search engines, while Table 2 summarizes the main topics of the top 10 hits returned by both search engines. Searching for another keyword, in this case "senate", we now expect to retrieve blogs that have information regarding the appointments to the Canadian Senate that occurred in March 2005. BlogCrawler again provides the results we expect as Table 3 demonstrates. Hence, we can conclude that by feeding a specific type of site into the web crawler as the root node in the social network, and following the links to other blogs, we are able to restrict the search engine to a specific topic without needing to enter clarifying keywords in the search query.

Table 2: Topics of Top 10 Hits for "budget" Keyword Search

BlogCrawler		Technorati	
Blog	**Topic**	**Blog**	**Topic**
Warren Kinsella	Provincial budget	Childe Roland	Film
CalgaryGrit	Federal budget	Smoke Filled Doom	Israel budget
Monte Solberg	Federal budget	Deficient Brain	American budget
Globe and Mail	Federal budget (False positive)	Brewtown Politico	State budget
My Blahg	Federal budget	Brewtown Politico	American budget
Living in a Society	Federal budget	Mac Professionell	(German language post)
TDH Strategies	Federal budget	Oleg Dulin	American budget
On the Fence	American budget	unLively Lives	Personal diary
Freethought.ca	International trade	The Reminisce Mind	Personal diary
Capitalist Pig vs. Socialist Swine	Federal budget	New Leadership Blog	American election

BlogCrawler search performed for "budget" on 2005/03/31 at 2:45pm. Technorati search for "budget" performed on 2005/03/31 at 2:47pm.
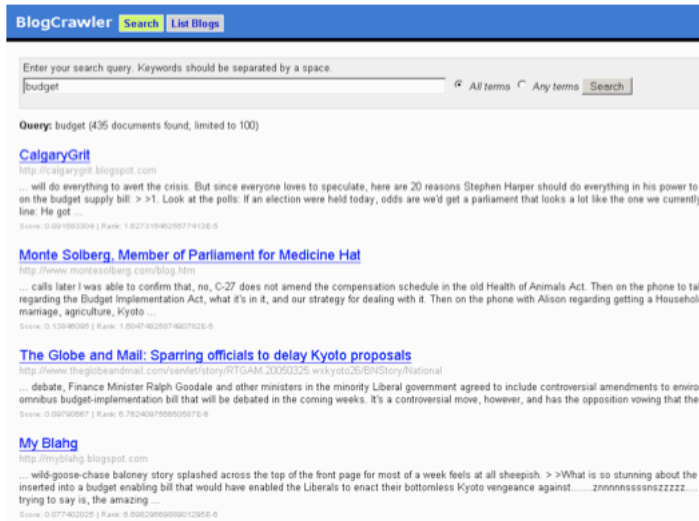
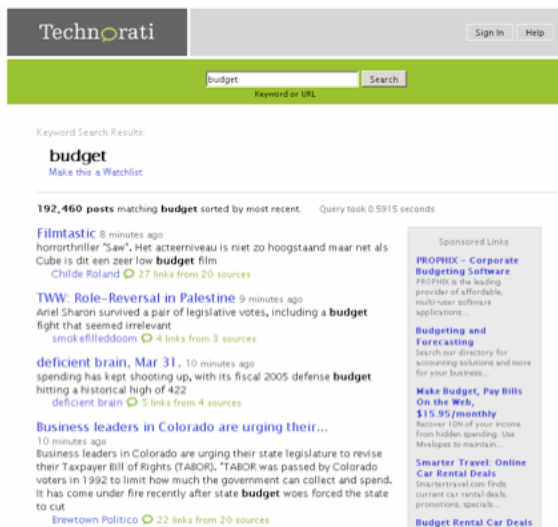**Figure 7: Results Returned for BlogCrawler**

Figure 8: Results Returned for Technorati

We also discussed the idea that blogs returned from searches should also be trustworthy. From our analysis of social networks, we concluded that if we rank blogs based on the number of links to it, then we should be able to order our collection of blogs by relative importance within the greater blogosphere. Since our search engine sorts hits by the rank of a blog, this leaves us to show that the actual page ranks calculated by BlogCrawler are accurate. Qualitatively, we see that the rankings calculated by BlogCrawler, supplemented by the biasing discussed in Section 4.5 appear to reflect this idea. Indeed, the sites we deemed "expert" blogs when calculating the page ranks all appear at or near the top of the ranked list of blogs contained in the index. We also note that the problem of the BlogValidator incorrectly identifying sites as blogs is overcome by the fact that those sites tend to have low page ranks relative to actual blogs.

Table 3: Topics of Top 10 Hits for "senate" Keyword Search

BlogCrawler		Technorati	
Blog	Topic	Blog	Topic
CalgaryGrit	Canadian Senate	Semidi	American Senate
Peace, order, and good government	American Senate	Swing State Project	American election
My Blahg	Canadian Senate	Gluetree.com	Star Wars trailer
On the Fence	Canadian Senate	Dudes Drivel	American Senate
Capitalist Pig vs. Socialist Swine	Canadian Senate	Dudes Drivel	American Senate
Highway 401 Blog	Canadian Senate	Centerfield	American Senate
Warren Kinsella	Government audit	Uncountable Spoons	American Senate
Lotusland	American Senate	Blast Off!	Obituary
BlogsCanada	Canadian Senate	Citrus Commando	American election
Crawl Across the Ocean	Canadian Senate	LANL: The Real Story	University/student politics

BlogCrawler search performed for "senate" on 2005/03/31 at 3:38pm. Technorati search for "senate" performed on 2005/03/31 at 3:38pm.

5.3 Overall Results

Based on the results we have observed, we can safely conclude that BlogCrawler is successful in implementing a solution demonstrating the properties of the blogosphere and the power that social networks have. This is not to say that the system is ideal, as will be discussed in the following section. However as a prototype, BlogCrawler was able to provide us with quite accurate and very relevant hits on our generalized keyword queries. The system was also able to rank blogs according to our personal preferences as defined by the blogs we used to bias the page rank calculations. As a basis for evaluating the use of social networks within the application of searching, the algorithms succeeded.

6. CONCLUSIONS

As we noted in the opening of this paper, blogs are increasingly being used by a variety of users in the field of media analysis. Blogs inherently express the thoughts of the public at large, so the importance of knowing what is being said on the many blogs on the Internet is important. With this comes the requirement to be able to search for what types of views we want in a more efficient fashion than is usually used. Indeed, unlike searching for a particular piece of information, looking for a specific string of text, media analysis requires people to look for a specific topic, but any type of viewpoint. Searches like this require the use of more general keyword queries, which we showed earlier to be cumbersome in terms of the amount of irrelevant sites the user would have to filter out. Therefore, this paper ultimately needed to answer two questions. First, do blogs exhibit properties of a social network? Secondly, can we use these properties to provide a level of human intelligence to search results while keeping the actual searching process simple? With BlogCrawler successfully implementing a solution that satisfies our requirements we can conclude several things regarding blogs in particular and social networks in general.

First, we can conclude that the blogosphere does constitute a social network with its interconnected set of hyperlinks and references. As our web crawler demonstrated, given any blog, its set of outgoing links linked it to other blogs of the same type, exhibiting the property of clustering, forming a community structure of blogs. The cluster of blogs we crawled all tended to be focus on the same issues, even if the opinions and viewpoints expressed on them differed. Moreover, we observed that the set of hyperlinks from one blog to another exposed embedded trust relationships in the blogosphere, with the most important and trustworthy blogs being linked to the most. Blogs clearly do constitute a social network.

From our experimental search results, we also can conclude that using social networks in a searching application can make the search process much more efficient. By limiting the search engine to a specific cluster of blogs, we see that effectively filter out irrelevant blogs that we would normally have had to filter out using contextualizing keywords. This also limits the numbers of blogs to only those with some level of authority since blogs that no one reads or trust will hardly ever be linked to. Social networks, then, add a modicum of human intelligence into the search process by recognizing and leveraging the human efforts made when constructing the social network in the form of linking to other blogs.

Interestingly, the conclusions we have found through our analysis of the blog social network also means that communities of blogs can be hierarchical. Blogs that focus on Jazz music form their own social networks, but may fall under a more general category of music blogs. Indeed, we would go so far as to say that the entire blog community may be hierarchical – using ontologies and intelligent agents to analyze and categorize individual blogs it is entirely possible to generate an overall blog network topology [6]. We have already seen in a simple manner how we can use the inherent organization of social networks to aid us in implementing practical applications for common problems. Knowledge of how the entire blogosphere is structured would, therefore, bring many benefits. We also believe that ideas from Referral Networks [20] can further enhance the properties of the system described here by dynamically restructuring the neighbourhood of blogs used in search process.

In all facets of life, we use social networks because we recognize that the relationships we build with others are valuable and useful. In the world of Computer Science, social networks allow us to add a human element to networking problems by recognizing the inherent organization structure social networks provide and realizing that there is information imparted whenever we decide to connect one resource to another. The common problem with searching is often that our search engine is not intelligent enough to recognize what we want. As we have seen with the BlogCrawler project, by using social networks, search engines can be intelligent enough to ensure that we are always satisfied by what we get back.

7. REFERENCES

[1] Blom, Thom et. al. (1998) *Writing a Web Crawler in the Java Programming Language*. Retrieved 12 September 2009 http://java.sun.com/developer/technicalArticles/ThirdParty/WebCrawler/.

[2] Eiron, Nadav et. al. (2004) "Ranking the web frontier," *Proceedings of the 13th international conference on World Wide Web*, pp. 309-318.

[3] Girvan, M. and Newman, M.E.J. (2002) "Community structure in social and biological networks," *Proceedings of the National Academy of Sciences*, Volume 99, Issue 12, pp. 7821-7826.

[4] Google (2009) *Google*. Retrieved 30 October 2009 from http://www.google.com.

[5] Haveliwala, Taher H. (2002) "Topic-Sensitive PageRank," *Proceedings of the 11th World Wide Web conference*, pp. 517-526.

[6] Heflin, Jeff (2004) *OWL Web Ontology Language Use Cases and Requirements*. Retrieved 4 April 2009 from http://www.w3.org/TR/webont-req/.

[7] Henzinger, Monika. (2000) "Link Analysis in Web Information Retrieval," Bulletin of the IEEE Computer Society Technical Committee on Data Engineering.

[8] Herring, Susan C. et. al. (2004) "Bridging the Gap: A Genre Analysis of Weblogs," *Proceedings of the 37th Hawaii International Conference on System Sciences.*

[9] Herring, Susan C. et. al. (2005) "Conversations in the Blogosphere: Analysis 'From the Bottom Up'," *Proceedings of the 38th Hawaii International Conference on System Sciences*.

[10] Kleinberg, Jon (1999) "The Small-World Phenomenon: An Algorithmic Perspective," *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, pp. 163- 170.

[11] Krishnan, Sriram (2004) *Writing a web crawler*. Retrieved 15 November 2009 from http://dotnetjunkies.com/WebLog/sriram/archive/2004/10/10/28253.aspx.

[12] Lindahl, Charlie & Blount, Elise. (2003) "Weblogs: Simplifying Web Publishing," *Computer*, Volume 36, Issue 11, pp. 114 -116.

[13] Nanno, Tomoyuki et. al. (2004) "Automatically Collecting, Monitoring, and Mining Japanese Weblogs," *Proceedings of the 13th international World Wide Web conference*, pp. 320-321.

[14] Nardi, Bonnie A. et. al. (2004) "Blogging as Social Activity, or, Would You Let 900 Million People Read Your Diary?" *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, pp. 222-231.

[15] Newman, M.E.J. (2001) "The structure of scientific collaboration networks," *Proceedings of the National Academy of Sciences*, Volume 98, Issue 2, pp. 404-409.

[16] Rodrigues, Maira Ribeiro et. al. (2003) "A System of Exchange Values to Support Social Interactions in Artificial Societies," *Proceedings of the Second Annual International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 81-88.

[17] Slashdot (2009) *Slashdot*. Retrieved 13 December 2009 from http://slashdot.org.

[18] Technorati (2009) *Technorati*. Retrieved 14 December from http://www.technorati.com.

[19] Wellman, Barry (1996) "A Sociological Perspective on Collaborative Work and Virtual Community," *Proceedings of the 1996 ACM SIGCPR/SIGMIS Conference*, pp. 1-11.

[20] White, T., McQuaker, S., Salehi-Abari, A., (2008) "On the importance of relational concept knowledge in referral networks. Artificial Intelligence Review Vol. 29, Issue 3-4 pp. 287-303.

[21] Kleinberg, J.M. (1999) "Authoritative sources in a hyperlinked environment", Journal of the ACM, 46(5) pp. 604-632.