

PAC'nPost: A Framework for a Micro-Blogging Social Network in an Unstructured P2P Network

H. Asthana
Department of Computer Science
University College London
Gower St., London WC1E 6BT, UK
h.asthana@cs.ucl.ac.uk

Ingemar J. Cox
Department of Computer Science
University College London
Gower St., London WC1E 6BT, UK
ingemar@cs.ucl.ac.uk

ABSTRACT

We describe a framework for a micro-blogging social network implemented in an unstructured peer-to-peer network. A micro-blogging social network must provide capabilities for users to (i) publish, (ii) follow and (iii) search. Our retrieval mechanism is based on a probably approximately correct (PAC) search architecture in which a query is sent to a fixed number of nodes in the network. In PAC, the probability of attaining a particular accuracy is a function of the number of nodes queried (fixed) and the replication rate of documents (micro-blog). Publishing a micro-blog then becomes a matter of replicating the micro-blog to the required number of random nodes without any central coordination. To solve this, we use techniques from the field of rumour spreading (gossip protocols) to propagate new documents. Our document spreading algorithm is designed such that a document has a very high probability of being copied to only the required number of nodes. Results from simulations performed on networks of 10,000, 100,000 and 500,000 nodes verify our mathematical models. The framework is also applicable for indexing dynamic web pages in a distributed search engine or for a system which indexes newly created BitTorrents in a de-centralized environment.

General Terms

Dissemination, Retrieval

Keywords

Microblogging, Social Network, Unstructured P2P Network, Restrained and Rapid Dissemination, Probabilistic Retrieval

1. INTRODUCTION AND PRIOR WORK

P2P networks can be generally categorized into two classes, namely structured and unstructured networks. Structured networks, typically based on distributed hash tables (DHTs), bind data to designated locations within the network. However this binding makes structured networks particularly susceptible to adversarial attack [6], and to churn, i.e. the random entry and exit of nodes to the network. Unstructured networks exhibit no such binding between data and nodes. As such, they are much less affected by churn, and are generally more resistant to adversarial attack. However, since a particular document being sought by a user can be anywhere in the network, search in an unstructured P2P network is necessarily probabilistic.

We assume search of an unstructured network is based on a probabilistic search that queries a fixed number of nodes. Copyright is held by the author/owner(s).

WWW 2012 Companion, April 16–20, 2012, Lyon, France.
ACM 978-1-4503-1230-1/12/04.

This form of search is known as Probably Approximately Correct (PAC) search [3, 4, 1], and a strong theoretical framework exists to predict the probability of a successful search based on the distribution of documents in the network. Given a set of n nodes in the network, we assume that the object of interest is stored on a random subset of r nodes. A query is issued to a random subset of z nodes. We are interested in the probability that the two subsets have a non-empty intersection, as this implies a successful search for that object. The overall expected retrieval accuracy α is

$$\alpha = 1 - \left(1 - \frac{r}{n}\right)^z \quad (1)$$

where α is the proportion of blogs¹ retrieved by querying z nodes relative to performing an exhaustive search.

Rumour spreading algorithms, also known as gossip spreading protocols or epidemic protocols, provide an efficient way to rapidly spread information within a network. One of the earliest uses of rumour spreading was by Demers *et al.* [5] to synchronize replicated databases. Apart from spreading a document across a network, rumour spreading algorithms can also be used for membership and random sampling [2]. Previous work on rumour spreading usually has the goal of ensuring that the rumour spreads to the entire network at the fastest possible rate. Our requirement is subtly different; we require that a document be replicated rapidly only onto a fraction of the network, thereby guaranteeing that the replication rate is sufficient for a PAC search to have a high probability of success.

2. THE PAC'NPOST FRAMEWORK

Our goals are two-fold. First, a user should be able retrieve the blogs of other users he/she is *following*, with a sufficiently high accuracy by querying z random nodes in the network. Second, a user should be able to perform a keyword search, by sending the query to, again, z random nodes. To retrieve the required blogs, a node in the network makes a *blog-request* to z other nodes every s seconds. We assume that the interval between requests, s , is 30 seconds, which allows for a user experience that is similar to Twitter.

We assume that the required accuracy is 95% and that the number of nodes queried is fixed at 25. An accuracy of 95% implies that on average we will retrieve 95% of the blogs we want by querying z nodes, as compared to searching the entire network. Both assumptions are arbitrary. The decision to query 25 nodes reflects the need to minimize network bandwidth and latency for the user. Setting $\alpha = 0.95$ and $z = 25$, we get $\frac{r}{n} = 0.12$, i.e. a blog needs to be

¹Henceforth, *tweet* and *blog* are used interchangeably with *micro-blog*

Table 1: Information retrieval and dissemination with a creation rate of 25 blogs per second for various network sizes - μ_a is the average retrieval accuracy, μ_r is the average replication rate (% of network), and σ denotes the corresponding standard deviation.

$n = 10,000$					$n = 100,000$					$n = 500,000$				
$ \mathcal{T} $	Retrieval		Dissemination		$ \mathcal{T} $	Retrieval		Dissemination		$ \mathcal{T} $	Retrieval		Dissemination	
	μ_a	σ_a	μ_r	σ_r		μ_a	σ_a	μ_r	σ_r		μ_a	σ_a	μ_r	σ_r
1	90.652	0.421	7.371	3.752	2	93.509	0.180	11.807	6.518	2	92.044	0.312	11.647	6.652
2	97.045	0.177	13.055	6.236	3	96.676	0.147	16.635	8.699	3	95.657	0.154	16.446	8.903
3	98.701	0.114	18.058	8.182	4	98.074	0.130	20.947	10.468	4	97.383	0.223	20.700	10.606

replicated to 12% of the network. The accuracy measure applies to both the retrieval of the *followed* blogs as well as to searching the contents of all the blogs in the network.

Next, we assume that each node in the network contributes, on average, 1GB of disk space to support the services. Of this contributed disk space, 90% is utilized for storage of the blogs, and the remaining 10% for indexing the stored blogs.

When a node queries other nodes, or is queried by another node, it transfers its most recent blog, thereby spreading the blog into the network. To increase the rate of spreading, we introduce a *transfer buffer*, \mathcal{T} , which stores a small fraction of the most recent blogs that the peer has encountered. The transfer buffer has a fixed size, \mathcal{T}_L . To retrieve the blogs Node \mathcal{A} is *following*, it makes a *blog-request* to z other nodes every s seconds. A request consists of 1) the list, \mathcal{L}_A , of the user ID's Node \mathcal{A} follows 2) the most recent blog created at Node \mathcal{A} , and 3) the blogs contained in the *transfer buffer* \mathcal{T}_A of Node \mathcal{A} . When Node \mathcal{B} receives a *blog-request*, it sends back a *blog-response*, which consists of 1) the most recent blogs of all the peers specified in \mathcal{L}_A , which are found in its storage area, 2) the most recent blog created at Node \mathcal{B} , and 3) the blogs contained in its *transfer buffer* \mathcal{T}_B .

To perform a keyword search, the list \mathcal{L} is replaced by the keyword query in the *request*, and the *response* contains the blogs matching the query, which are then merged and re-ranked at the originating node.

At the completion of a request/response, the node stores all the new blogs it has come across into its storage area. It then reconstructs its transfer buffer. A number of strategies for selecting the blogs for the buffer are possible. We investigated two, namely most recent, and a probabilistic strategy in which the probability of selection is proportional to the age of the blog. The performance of the selection methods was nearly identical with the deterministic method performing marginally better. We omit the results for the probabilistic methods due to space limitations. The goal of selecting blogs for the transfer buffer is to continue spreading newer blogs at the expense of older blogs.

3. SIMULATIONS

To verify our theoretical framework, we performed simulations based upon the parameter values discussed in the previous section. At each iteration, 25 new blogs are created at random nodes in the network. This is the normalized rate of blog creation for a network of 1M nodes based upon Twitter's average of 2,300 tweets per second and a 100 million *active* user base. The simulations are based on network sizes of 10,000, 100,000, and 500,000 nodes. Since z is small, we assume that each iteration is one second in magnitude, i.e. the latency is less than 1 second, and that iteration and second can be used interchangeably. Under steady-state conditions, i.e. when all transfer buffers are full, we recorded each blog's replication rate at the end of each simulation. Retrieval accuracy was also measured. We assumed each node follows 10 randomly chosen user ID's.

Table 1 summarizes the results. For blogging rates of 25 per second, we observe that a buffer size of 2 or 3 is sufficient to provide accuracies of 95% or higher. The average replication rate was slightly higher than the required 12%, but did not exceed 17%. The mean and standard deviations are the average over the 10 trials, for each configuration.

4. CONCLUSIONS AND FUTURE WORK

This paper considered the design of a micro-blogging social network in an unstructured peer-to-peer network. It consists of a PAC search together with a constrained rumour spreading algorithm. Simulations showed that the rumour spreading algorithm successfully spreads blogs to a fraction of the network (about 12%), which is sufficient to guarantee that the retrieval accuracy is at least 95%. This was achieved with a small transfer buffer size of just 2 or 3.

Our current on-going work focuses on churn. Since unstructured networks do not have to maintain a DHT, the effect of churn (i.e. nodes joining and leaving) on the proposed framework would be to increase the number of nodes queried or the replication rate to achieve the same retrieval accuracy. We aim to incorporate the expected levels of churn as well as massive leave/joins into future simulations.

Analysis of Twitter shows that the blog creation rate of individual users is not constant, but follows a power law distribution. Similarly, the *follower* and *followed by* statistics also follow a power law distribution. Furthermore, the blog creation rate contains spikes where the rate increases five fold when an important event occurs. We plan to incorporate these statistics into our framework and use them to arrive at a system-wide accuracy metric as part of our future work. From our experiments, it is clear that the required size of the transfer buffer depends on the network size and the blog creation rate. We intend to investigate an *adaptive* transfer buffer strategy where a node would select the size of the transfer buffer based upon its local knowledge of the network size, blog creation rate, and an estimate of the replication rate of blogs.

5. REFERENCES

- [1] ASTHANA, H., FU, R., AND COX, I. J. On the feasibility of unstructured peer-to-peer information retrieval. In *ICTIR* (2011).
- [2] BORTNIKOV, E., GUREVICH, M., KEIDAR, I., KLIOT, G., AND SHRAER, A. Brahms: Byzantine resilient random membership sampling. *Computer Networks* 53, 13 (2009), 2340–2359.
- [3] COX, I. J., FU, R., AND HANSEN, L. K. Probably approximately correct search. In *ICTIR* (2009).
- [4] COX, I. J., ZHU, J., FU, R., AND HANSEN, L. K. Improving query correctness using centralized probably approximately correct (pac) search. In *ECIR* (2010).
- [5] DEMERS, A., GREENE, D., HAUSER, C., IRISH, W., LARSON, J., SHENKER, S., AND STURGIS, H. Epidemic algorithms for replicated database maintenance. In *Proceedings of the 6th annual ACM Symposium on Principles of distributed computing* (1987).
- [6] URDANETA, G., PIERRE, G., AND STEEN, M. A survey of dht security techniques. *ACM Computing Surveys (CSUR)* 43, 2 (2011), 8.