

Evaluation of Models for Analyzing Unguided Search in Unstructured P2P Networks

Bin Wu and Ajay D. Kshemkalyani

Computer Science Department, Univ. of Illinois at Chicago, Chicago, IL 60607, USA
{bwu, ajayk}@cs.uic.edu

Abstract. Evaluating the efficiency of unguided search based on random walk in unstructured peer-to-peer networks is important because it provides guidelines in correctly setting the parameters of the search. Most existing work is based on simulations. We evaluate two analytical models – the algebraic model and the combinatorial model – for various search efficiency metrics against simulation results. We use the random graph topology and assume unguided searches. The results show that the two analytical models are accurate and match each other closely. We study the impact of the average node degree, hop count, number of walkers, and replication ratios on node coverage, object recall, and message efficiency, and on the accuracy of the models.

1 Introduction

Mechanisms for indexing and searching for objects are at the heart of peer-to-peer networks [9, 12]. Although current research tends to favor distributed but structured architectures over unstructured ones, there are several drawbacks of structured architectures [12]: they are not suited for systems with very rapid churn, they incur overhead for maintaining the logical structure, and they do not support keyword searches or complex queries such as range queries. Unstructured overlays (e.g., Gnutella) have been more widely used to build P2P applications. Further, in ubiquitous and ad-hoc mobile networks, the requirement for a structure imposes additional organizational restrictions. Thus, unstructured overlays are also favored in such environments.

Recent research on unstructured P2P networks has focused on the search strategies and replication schemes ([3], [12]). The goal is to avoid message explosion, achieve good performance, and maintain the simple and flexible topology of an unstructured P2P overlay. The improvements over the traditional flooding are the expanding ring flooding and random walk [10]. These approaches can be classified as *unguided searches*. In contrast, *guided searches* remember some specific information on the network topology or on the past searches. When forwarding a query message in search of an object, this information is used to narrow down the choice of the neighbor(s) to forward the query to [4], [15].

In unstructured overlays, an object is not usually identified by its unique ID, and it may not even have an object ID. Thus, query by keyword is a primary

method of indexing and searching for objects in such environments ([11],[8]). For keyword searches, the “matching” depends on the relevancy between an object and the set of keywords used in the query. Keyword based search methods are usually closely related to semantic identification and information retrieval techniques ([18],[13],[14]) such that an effective semantic based clustering and a pertinent model on the characteristics and distributions of keywords ([1]) are essential.

We focus on unguided searches in unstructured networks. Most existing work on performance is based on empirical and simulation studies. The most comprehensive analysis of search and replication strategies is by Lv, Cao, Cohen, Li, and Shenker [10]. They studied the impact of network topology, query distribution, and replication distribution using modeling and simulations. Based on models from queuing theory, [6] presented the performance of different indexing approaches, in terms of system throughput and probability of “successful query”. Using Chernoff bounds on a well-connected graph, Gkantsidis et al. [7] showed that the effect of a k -step random walk is statistically similar to that of taking k independent samples. Using this result, expressions for the success rate, message overhead, and time overhead of random walk as functions of the Time-to-Live (TTL), object popularity, and number of walkers were given [2].

Overview of Results [R1.] We evaluate the goodness of two analytical models – the algebraic model and the combinatorial model [16] – against simulation results, for various search efficiency metrics. We use the random graph topology and assume unguided searches. The search metrics we consider are the node coverage, the message efficiency, and the object recall. The results show that the two analytical models are very close to each other and reasonably accurate. Using simulations, we evaluate the impact of the parameters such as the average node degree, hop count, number of walkers, and replication ratios on node coverage, object recall, and message efficiency, and on the accuracy of the models. [R2.] We then suggest how to enhance the analytical models to also account for the average node degree in the random graph and improve accuracy.

The models provide valuable confirmation of the extensive simulation results of Lv et al. [10]. The models also help to understand how the settings of the various parameters impact the efficiency of the search strategies [16]. This allows system designers to tune the parameters to achieve performance trade-offs.

2 Assumptions and Background

Graphs. The simulations from [10] indicate that random graphs provides better performance than Power-Law Random Graphs and Gnutella graphs, and they also have a much more uniform degree distribution. So we assume a random graph topology which is also a representative small-world model. Between any two nodes in a random graph, there is a link with a probability of p . Let the graph have N nodes. The expected number of links is $p\frac{N(N-1)}{2}$. The average node degree, denoted D , is $p(N-1)$. Random graphs have a rich set of properties [5].

Metrics. Message overhead (number of messages), time overhead (number of hops), and the success rate are typically used to evaluate searches. We showed that *node coverage*, the fraction of nodes that have been visited, is a very useful metric because of the following [16].

1. It is useful to estimate success rate.
2. For keyword and range searches, it is useful to measure *object recall* and *message efficiency*.
 - *Object recall* is the number of objects found that satisfy the query criteria.
 - *Message efficiency* is the object recall per query message.
3. It is useful to determine when to call off a search.
4. It is useful to estimate the replication ratio of an object.

We assume a *stateless* search, i.e., nodes do not have any memory of previous searches or forwardings. Two analytical models to calculate node coverage, and hence object recall and message efficiency, were given in [16].

2.1 The Algebraic Model

In the random graph, a message forwarding may visit a node that was already discovered. Let x denote the number of query messages and $u(x)$ be the number of distinct nodes visited by those messages so far. Then

$$u(x+1) = u(x) + \frac{N-u(x)}{N-2} \quad (1)$$

This equation can be solved as:

$$u(x) = C e^{-\frac{x}{N-2}} + N \quad (2)$$

Assuming W random walkers, the solution for equation (2) is:

$$u(x) = \begin{cases} N - (N - W - 1)e^{-\frac{W+1-x}{N-2}} & \text{for } x > W + 1 \\ x & \text{for } x \leq W + 1 \end{cases} \quad (3)$$

Let H be the number of query hops and $x = WH + 1$. Then $u(H)$ is:

$$u(H) = \begin{cases} N - (N - W - 1)e^{-\frac{W(1-H)}{N-2}} & \text{for } H > 1 \\ W + 1 & \text{for } H = 1 \end{cases} \quad (4)$$

2.2 Combinatorial Model

Let v be the number of nodes visited so far. Let $Pr(u, v)$ denote the probability that after v node visits, u distinct nodes have been visited. Then, for $u \leq v$:

$$Pr(u, v)|_{u \leq v} = \begin{cases} 1 & \text{if } u = 2, v = 2 \\ Pr(u, v-1) \cdot \frac{u-2}{N-2} & \text{if } u = 2, v \neq 2 \\ Pr(u-1, v-1) \cdot \frac{N-(u-1)}{N-2} + Pr(u, v-1) \cdot \frac{u-2}{N-2} & \text{if } 2 < u < v \\ Pr(u-1, v-1) \cdot \frac{N-(u-1)}{N-2} & \text{if } u = v > 2 \end{cases} \quad (5)$$

Define the expected number of distinct nodes covered by W random walkers after traveling H hops to be $\bar{u}(H, W)$.

$$\bar{u}(H, 1) = \begin{cases} H + 1 & \text{if } H \leq 2 \\ \sum_{i=3}^{H+1} Pr(i, H + 1) \cdot i & \text{if } H > 2 \end{cases} \quad (6)$$

To compute $\bar{u}(H, W)$ for $W > 1$, assume they travel the network sequentially. A walker may see some nodes that have already been visited by earlier walkers. Let $new(H, i)$ denote the expected number of distinct new nodes visited by the i th walker. We can express $new(H, i)$ as follows:

$$new(H, i) = \left[1 - \frac{\sum_{k=1}^{i-1} new(H, k)}{N} \right] \bar{u}(H, 1) \quad (7)$$

The expected total number of distinct nodes visited by W walkers is:

$$\bar{u}(H, W) = \sum_{i=1}^W new(H, i) \quad (8)$$

3 Experimental Evaluation of Models

To test the validity of the algebraic and combinatorial models, we simulated random walk on a undirected random graph having $N = 20,000$ nodes. The graph and search parameters included: N , D , W , H , and R (the replication factor). We measured the node coverage, object recall, and message efficiency. Both models performed very close to each other; hence in the graphs for most experiments, we show their values as a single plot. A study of the small differences between the models is shown in Section 3.4.

3.1 Node Coverage

We compare node coverage as computed from our analytic models and that obtained from the simulations. We study the impact of D and W on node coverage.

Effect of node degree. We simulated 32 random walkers and varied the average node degree. From Fig. 1(a), observe that both models give higher values of node coverage than the simulation results. When the node degree is small ($D=4$, $D=6$) the difference between analytical results and simulations is large (up to over 30% for $H=600$) but this difference reduces as D increases. For $D=20$, the difference remains below 3% at all values of H (see Fig. 1(b)).

The analysis models consider each “next step” of message forwarding as random and the probability of visiting a new node is determined by the current node coverage only; the node degree is not taken into account. Consider a node i that is being visited for the second time. For the forwarding to be done by i , ideally i must select a neighbor that was not visited before. However, as we assume a




Fig. 1. Impact of average node degree on node coverage. (a) Absolute values. (b) Relative differences with respect to models.




Fig. 2. Impact of number of walkers on node coverage. (a) Absolute values. (b) Relative differences with respect to models.

stateless search, i may forward the message along an already traversed link. This reduces the probability of forwarding to an unvisited node, but the models do not account for this effect. So the coverage computed from our models is always higher than simulation results. However, this effect is reduced as average node degree increases because for a higher degree node, the chance that a subsequent forwarding is along an untraversed link is higher than that for a low degree node. The simulation results should match the model as node degree goes to N .

Observe from Fig 1(b) that for all values of node degree, the differences between analytical and simulation results increase as the hop count H increases. This is because as messages and hops increase, the fraction of nodes that have been visited before also increases, magnifying the effect of node degree in reality.

Effect of number of walkers. Figure 2(a) shows the comparison of node coverage for 32, 64, and 128 walkers in the random graph with $D = 10$. As observed above, the simulation results are consistently smaller than our analytical results and the difference increases as the message number increases. Our simulations show no significant impact of the number of walkers on node coverage as a

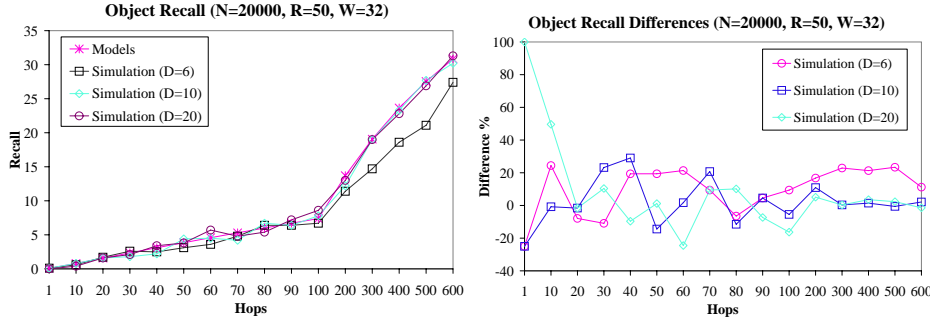


Fig. 3. Impact of average node degree on object recall. (a) Absolute values. (b) Relative differences with respect to models.

function of message overhead. Fig 2(b) shows the relative deviation of each of the simulation cases from analytical results. The fluctuations observed are likely due to our limitation of small sample space for the random sampling.

3.2 Object Recall

A query can be considered as more efficient if a certain number of query messages yields higher object recall. In the following simulations, we inspect the influence of average node degree and replication ratio, respectively.

Effect of node degree. We simulated 32 random walkers searching for an object with 50 randomly distributed replicas. The average node degree was varied as 6, 10, and 20. In each case, we took the average of the object recalls for 10 searches. The simulation results are compared with the analytical models in Fig 3.

Note that the object recall for the analytical models are derived from node coverage. Since the actual recall value for a single search (run) heavily depends on the random choices made by each walker at each step, we expect certain fluctuation on the curves for the simulation results. (Ideally, the number of samples should be large enough to ensure stable results). Fig 3(a) indicates that the analytical models generate higher recall values than the simulations in general. The deviations of the simulation results from analytical models diminish as node degree increases. This is reasonable because the object recall is expected to increase as node coverage increases. Fig 3(b) shows the relative deviation from the analytical results. For a smaller number of hops, the object recall value obtained from simulations is too small to generate meaningful comparisons. As the recall increases with hop number, the effect of node degree becomes apparent – the higher the degree, the less the deviation from analytical models.

Effect of replication. We simulated a search in a graph with $D=10$ and $W=32$, while setting the replication ratio of the queried objects to 50 and to 200. The results are plotted in Fig 4(a). The analytical values are somewhat similar to (but




Fig. 4. Impact of degree of replication on object recall. (a) Absolute values. (b) Relative differences with respect to models.

a little greater than) the simulation results (Fig 4(b)), barring some exceptions that are likely due to the limitation of sampling spaces. With both replication values, the relative deviation tends to diminish as the hop number increases. This is probably because with more hops, the larger recall values recorded from each run produce more stable output than for the case of fewer hops.

3.3 Message Efficiency

Message efficiency is a derived quantity: (object recall) / (message overhead). The expected object recall is in proportion to node coverage. According to Equation (3), the efficiency should decrease as the message overhead increases. Since our analytic models provide an upper bound for expected node coverage, we also expect that the models give an upper bound for expected message efficiency. What interests us is to investigate how close the simulation results would approach the “expected” upper bound, and what are the effects of search and topology parameters upon this approximation. The results indicate that it is reasonable to use the analytical results from our models as upper bounds of the expected message efficiency.

Effect of number of walkers. The number of walkers is seen to have no impact on the message efficiency, as computed from the models. Fig. 5(a) compares the message efficiency obtained from our models and from simulation results for $W=32$, $W=64$, and $W=128$, with 50 replicas per object. In our test cases, when the message overhead is low, the recall values obtained from the simulations are still not stable enough for comparison. As the results smooth out with increasing messages, the analytical results tend to have better message efficiency than the simulation results. Also, the simulations suggest that the number of walkers has marginal impact on message efficiency. The relative differences between the analytical results and simulations are generally below 10% for all the 3 cases when message overhead > 6400 .




Fig. 5. Impact on message efficiency. (a) Impact of number of walkers. (b) Impact of average node degree.




Fig. 6. Node coverage for the algebraic and combinatorial models. (a) 32 walkers. (b) 128 walkers.

Effect of average node degree. The simulation results in Fig. 5(b) show a similar trend as Fig. 5(a) when we vary the average node degree while fixing the other parameters. As D has an impact on node coverage (see Section 3.1), the relative performance of message efficiency for different D tends to be the same as the relative performance of node coverage as D varies.

3.4 Comparison of the Analytical Models

The two analytical models fit each other extremely well. The node coverages generated from either model are so close to each other that we represented the results from both as a single plot in the graphs so far. Figure 6 shows the relative difference in terms of node coverage, when $D=10$ and W is 32 and 128, respectively. In both cases, the algebraic model generates slightly higher value of node coverage than the combinatorial model. Their differences are at most 0.3% for $W=32$ and at most 0.1% for $W=128$. Both models can serve as an upper bound for the estimate of node coverage and object recall.




Fig. 7. Refining the analytical model to account for node degree.

4 Discussion

Our simulations indicate that the algebraic and combinatorial formulas provide an upper bound on the node coverage, and the actual value of node coverage approaches the upper bound as the average node degree increases (tending to N). On observing the analytical expressions, we see that the degree is not a parameter in the expressions for node coverage, and implicitly it is assumed that $D = N$. To account for the average degree D in the stateless routing models, we formulate the following approximation.

$$u(x+1) = u(x) + \frac{u(x)}{N} \cdot \left[\frac{D-d}{D} \right] \cdot \frac{N-u(x)}{N-2} + \frac{N-u(x)}{N} \cdot 1 \cdot \frac{N-u(x)}{N-2} \quad (9)$$

$\frac{u(x)}{N}$ is the probability that the current node being probed (node B in Fig. 7) has been visited before. The term $\left[\frac{D-d}{D} \right] \cdot \frac{N-u(x)}{N-2}$ gives the probability that the $x+1$ th forwarding visits a new node if the current node has been visited before. The term $1 \cdot \frac{N-u(x)}{N-2}$ represents this probability if the current node has not been visited before. d is the expected number of links of the current node B that have been traversed previously (across possibly multiple previous visits), given the current node coverage $u(x)$. Hence, in the second term, $\frac{D-d}{D}$ is the probability that the link randomly chosen by B for forwarding has not been explored before. Estimating d and the accuracy of the resulting models is studied in [17].

5 Conclusions

This paper evaluated the algebraic and the combinatorial models for computing node coverage (NC), object recall (OR), and message efficiency (ME) of unguided searches in random graphs. The results show that the two analytical models are accurate and match each other closely. The impact of the average node degree, hop count, number of walkers, and replication ratios on the accuracy of the models, as studied via simulations, is summarized.

1. The models give a little higher value of NC, OR, and ME than simulations.
2. As D increases, accuracy of models for NC increases.
3. As H increases (i.e., number of messages increases), accuracy of models for NC decreases.
4. The NC as a function of number of messages appears independent of W .
5. As D increases, accuracy of models for OR increases.

6. Increasing R and H increases the stability of OR values obtained from simulations, and increases the approximity of analytical to simulation results.
7. ME is seen to be almost independent of W .
8. ME from the simulations becomes more stable as H increases.
9. As D increases, the ME increases and accuracy of models also increases.

We are (i) extending our models to power law random graphs and other small world networks; (ii) enhancing our models to analyze other querying approaches, eg., guided searches, such as keyword-based and probability-based searches.

References

1. M. Bawa, G. S. Manku, P. Raghavan, SETS: Search Enhanced by Topic Segmentation. SIGIR 2003: 306-313
2. N. Bisnik, A. Abouzeid, Modeling and Analysis of Random Walker Search Algorithm in P2P Networks. Proceedings of HOT-P2P 2005
3. E. Cohen, S. Shenker, Replication Strategies in unstructured Peer-to-Peer Networks, ACM SIGCOMM, 177-190, 2002.
4. A. Crespo, H. Garcia-Molina, Routing Indices For Peer-to-Peer Systems. ICDCS 2002: 23-
5. P. Erdos, A. Renyi, Random Graphs. Publ. math. (Debrecen), Vol 6, p. 290-, 1959
6. Z. Ge, D. R. Figueiredo, S. Jaiswal, J. Kurose, D. Towsley, Modeling Peer-Peer File Sharing Systems. Proceedings of Infocom 2003
7. C. Gkantsidis, M. Mihail, A. Saberi, Random Walks in Peer-to-Peer Networks: Algorithms and Evaluation. Performance Evaluation, Volume 63: p. 241-263, 2006.
8. L. Liu, K. D. Ryu, K.-W. Lee, Supporting Efficient Keyword-based File Search in Peer-to-Peer File Share Systems. IEEE Global Internet and Next Generation Networks Symposium, Dallas, TX, November, 2004.
9. E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, S. Lim, A Survey and Comparison of Peer-to-Peer Overlay Network Schemes. IEEE Communications Survey and Tutorial, March 2004
10. Q. Lv, P. Cao, E. Cohen, K. Li, S. Shenker, Search and Replication in Unstructured Peer-to-Peer Networks. ICS 2002: 84-95
11. P. Reynolds, A. Vahdat, Efficient Peer-to-Peer Keyword Searching. Middleware 2003: 21-40
12. J. Risson, T. Moors, Survey of Research towards Robust Peer-to-Peer Networks: Search Methods. Computer Networks (to appear).
13. C. Schmitz: Self-organizing a Small World by Topic. P2PKM 2004
14. C. Tang, Z. Xu, S. Dwarkadas, Peer-to-peer Information Retrieval Using Self-organizing Semantic Overlay Networks. SIGCOMM 2003: 175-186
15. D. Tsoumakos, N. Roussopoulos, Adaptive Probabilistic Search for Peer-to-Peer Networks. Peer-to-Peer Computing 2003: 102-109
16. B. Wu, A.D. Kshemkalyani, Analysis Models for Blind Search in Unstructured Overlays, IEEE Network Computing and Applications, 2006.
17. B. Wu, A.D. Kshemkalyani, Predicting Unguided Search Metrics in Unstructured Overlays, UIC CS Technical Report, 2006.
18. D. Zeinalipour-Yazti, V. Kalogeraki, D. Gunopulos, On Constructing Internet-Scale P2P Information Retrieval Systems. DBISP2P 2004: 136-150