A Novel Probabilistic Flooding Strategy for Unstructured Peer-to-Peer Networks

Spiridoula V. Margariti
Department of Communication and Informatics
T.E.I. of Epirus
Arta, Greece, GR-47100
Email: smargari@cs.uoi.gr

Vassilios V. Dimakopoulos
Department of Computer Science
University of Ioannina
Ioannina, Greece, GR-45110
Email: dimako@cs.uoi.gr

***Abstract*—In this work we propose a novel probabilistic flooding strategy for unstructured p2p networks. Our strategy takes into account the popularity of resources and the hop distance from the node that initiated the query. The latter is used to estimate the number of nodes reached by the query message. Based on the above parameters we adjust the forwarding probability at the time a node receives the query message so as to reduce the duplicate message overhead while maintaining a high probability of query success. The primary goal of our approach is to minimize the cost of search associated with excessive message transmissions. The experimental results support our claims and provide insights into the effect of the above parameters to search performance.**

***Keywords*-distributed systems; duplicate messages; probabilistic flooding; unstructured p2p networks.**

I. INTRODUCTION

Unstructured peer-to-peer (p2p) networks consist of a large population of networked computers that offer resources and operate in a fully decentralized manner. Such systems are usually quite large and highly dynamic and each node (or computer) has only information for a small subset of the other participating nodes. A key challenge is that of locating a desired resource. In the absence of information about node and resource placement, these systems use mainly flooding and its variants to provide search facilities. In flooding, the node (peer) that initiates the search sends a query message to all its neighbors. Any neighbor that does not know about the resource propagates the message to all its neighbors, and so on, until the resource is discovered or some termination conditions are reached.

Although flooding is a fundamental strategy for resource discovery in several networks, it is responsible for overloading the network with a large number of messages [1]. Therefore, it is vital that search strategies in decentralized unstructured p2p systems be cost effective and successful. By cost effectiveness we mean that the produced traffic to spread a query over the network is within reasonable limits. A search is considered successful if it discovers at least one replica of the desired resource. Flooding-based search strategies are generally fast but produce excessive traffic, a large portion of which is redundant.

The objective of many works is to provide models and techniques that are capable to moderate the redundant traffic. A number of strategies, such as the ones based on random walkers [2], [3], [4], abandon flooding altogether. Others strive to limit the extend of flooding, usually in a probabilistic manner.

Probabilistic flooding, which is also the subject of this paper, is a class of search strategies which may exploit a property of the network to make a probabilistic decision whether a message should be forwarded to another node or not. More specifically, each node that receives a query message, forwards it to each of its neighbors with some probability p_f. This forwarding probability can be the same for all nodes [5], [6], [7], [8], [9] or varying [10], [11], [12]. In [5] a query message is forwarded to a fixed portion of neighbors (the authors used $p_f = 0.5$) while a value of $p_f = 0.6$ was chosen in [6] for exchanging routing information in power-law networks. Both [7] and [8] tune the value of p_f so that query messages reach all nodes with high probability, while [9] derives a relationship between the value of p_f and a desired node coverage level.

Adaptive Resource-based Probabilistic Search (ARPS) [11] utilizes different forwarding probabilities for different resources, depending on estimations of resource popularity. In addition, when applied to power-law networks each node adjusts this forwarding probability further, according to its degree, while in [10], the forwarding probability depends on both the sending and receiving nodes' degree. Finally, the authors in [12] study the effect of a forwarding probability which is decreasing exponentially with the distance from the node that posed the query.

In this paper we consider probabilistic flooding. In particular, our contribution is a novel strategy for varying the value of the forwarding probability based on the distance from the query initiator. However, in contrast to [12] the rate of decay is not exponential; it is based on an estimation of the node coverage, aiming at minimizing unnecessary message overhead. The estimation can be further enhanced by utilizing knowledge about resource popularity. All our results are backed by simulation experiments that evaluate the performance of the proposed strategy. In addition, we have conducted comparative measurements with other probabilistic flooding schemes, which demonstrate the superiority of our approach.

The rest of the paper is organized as follows. In the next section we summarize the system model and our basic assumptions. Section III presents our probabilistic flooding strategy and elaborates on the required calculations. In Section IV we give an evaluation of the proposed strategy through simulation experiments and finally, Section V concludes the paper.

II. MODEL AND ASSUMPTIONS

We consider an unstructured p2p network where peers (nodes) are organized in an undirected random graph. Let us assume a total of N nodes and an average node degree $\overline{d}$. Peers submit queries for locating a copy of a resource they are interested in. If there exist $r \geq 1$ replicas of a particular resource, then the *popularity* of the resource is defined as $q = r/N$, that is the ratio of peers that possess it.

It is a well known fact that when flooding is employed for locating a resource, a large portion of the generated traffic is redundant, i.e. the query message is re-transmitted, possibly multiple times, to already visited nodes. Due to this undesirable situation, *duplicate detection mechanisms* (DDMs) need to be employed in order to limit unnecessary duplicate messages. Such a mechanism is used for example in Gnutella where each query message is assigned a globally unique identifier (GUID) field [13]. When a peer receives a message, it stores its GUID in a local query cache and keeps it there for some time. If the peer receives the same query message again (i.e. the same GUID), it simply discards it, avoiding unnecessary retransmissions.

Of course, the mechanism can never be perfect in the sense that a duplicate message may always appear long after its GUID was removed from a peer's cache, for example due to delays in the underlying physical network. Nevertheless, this simple mechanism is quite powerful and eliminates most of the redundant traffic. Notice, however, that even by using a perfect DDM, there are certain duplicate messages that cannot be prevented. They concern copies of the query message that arrive at the same node through different paths in the network. All the DDM can do is stop propagating them *after* they arrive. In what follows, we assume that independently of the variation of flooding used, a perfect DDM is in effect.

III. PROPOSED STRATEGY

Consider a node u that has received a query for a particular resource. If u does not possess the resource, according to probabilistic flooding, it has to decide whether to forward it to a particular neighbor with some probability p_f. The main idea behind our strategy is that the query message should be propagated further *only if the query has not been answered yet*; that is p_f should be equal to the probability that the required resource has not been discovered by the flooding procedure up to that point. Thus p_f should be a function of t, the distance from the query initiator, or equivalently the "step" of the search.

Suppose that node u receives the query at step t, and that by step t there are $N(t)$ nodes that have received the query. Node u will decide whether to forward the message any further by estimating the probability that the resource has been found up to step t. The number of nodes that have not been visited yet is given by $N - N(t)$. If there exist r replicas of the required resource, then the probability that the resource has not been found as of step t is given by the probability that all those r replicas have been placed among the $N - N(t)$ remaining nodes (which clearly have to be $\geq r$ in number, otherwise the resource has been found already). Making the simplifying—but accurate as far as unstructured p2p systems are concerned, since there is no correlation among the topology and the placement of data—assumption of independent placements, the probability that a replica is placed in one of the not-yet-visited nodes is given by:

$$\frac{N - N(t)}{N} = 1 - \frac{N(t)}{N}.$$

Consequently, the probability that all r replicas are placed among these nodes is given by:

$$\left(1 - \frac{N(t)}{N}\right)^r.$$

This is the probability that none of the already visited nodes is an owner of the resource. In such a case, node u should indeed forward the query and this is exactly what we do in our algorithm.

Summarizing the above discussion, in our proposed strategy any node that receives the query message at step t, and does not possess the required resource, propagates it to its neighbors with a forwarding probability

$$p_f(t) = \left(1 - \frac{N(t)}{N}\right)^{qN}, \tag{1}$$

where $q = r/N$ is the popularity of the resource and $N(t)$ is the number of nodes that have received the query by step t. We let $p_f(0) = 1$, so that the initiator transmits the query to all its neighbors.

A. Estimation of the parameters

In order for a node to calculate the forwarding probability of (1), two quantities must be known: (a) the popularity of the resource (q) and the total number of visited nodes by step t ($N(t)$). Regarding (a), we assume that popularities are known. Otherwise they can be estimated at each node using techniques that monitor the local traffic (see for example [11]). In the absence of any knowledge or estimation, a peer may use a small value for q, to be on the conservative side.

We now turn to the estimation of $N(t)$. Let n_i be the expected number of nodes that lie in distance i from the query initiator Consider the message transmissions during step t. Because we have assumed that a duplicate detection mechanism is in effect, the only peers that forward the query during step t are the n_{t-1} nodes in distance $t-1$ from the initiator (otherwise, there would exist more than n_{t-1} nodes in that distance). Since a node should not transmit the message back to the neighbor that delivered it, the nodes in question will only transmit to new nodes in distance t. Given that the average node degree is $\overline{d}$, then from the $(\overline{d}-1)n_{t-1}$ message transmissions, only a portion of

$$1 - \frac{N(t-1)}{N}$$

will be delivered to new nodes. This is because, out of the N possible destinations of a messages, the $N(t-1)$ are already visited by time $t-1$. Consequently,

$$n_t = (\overline{d}-1)n_{t-1}\left(1 - \frac{N(t-1)}{N}\right)p_f(t-1), \qquad (2)$$

where we have accounted for the fact that nodes in distance $t-1$ forward the message with probability $p_f(t-1)$. Based on the above, we obtain:

$$N(t) = \sum_{i=0}^{t} n_i, \qquad (3)$$

where $N(0) = n_0 = 1$ (the initiator node). The recursive equations (1)–(3) can be used to obtain the desired forwarding probability.

IV. EVALUATION

In this section we evaluate the proposed strategy through simulation experiments. We have constructed a peer-to-peer message-level network simulator which is able to generate various topologies, such as random, random regular and power-law random graphs according to user-supplied parameters. After the topology construction, for each simulation run nodes of the system are selected in random and are marked as owners of a resource replica, according to given resource popularities. Then one of the peers, chosen uniformly randomly, initiates a search query. The search is limited by a TTL (time-to-live) parameter, t, which gives the maximum allowed number of steps / path length. During the search detailed statistics are kept, including the total number of messages, the number of duplicate messages, the number of visited nodes, etc. If at least one replica of the resource is found within the t steps, the query is considered *successful*, otherwise it is unsuccessful. We run each experiment at least 1000 times and average the results. The portion of runs that resulted in successful queries gives the *probability of success* for the particular value of t.

Table I
SIMULATION PARAMETERS (RANDOM NET, 100,000 NODES)

$\overline{d}$	p	flooding	mBFS	ARPS
4-regular	0.0001	$p_f = 1$	$p_f = 0.5$	$p_f = 0.9$
6	0.0003	$p_f = 1$	$p_f = 0.5$	$p_f = 0.8$

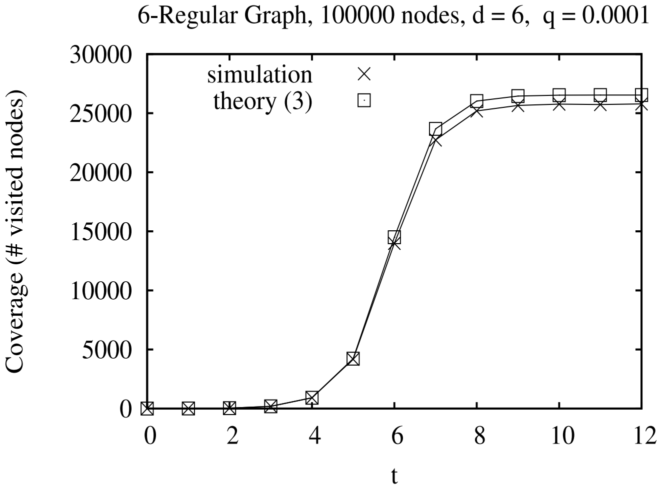

Figure 1. Coverage in a random 6-regular network of 100,000 nodes

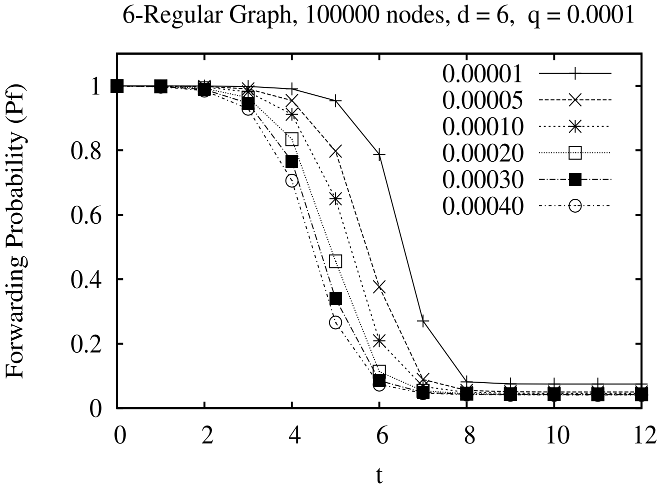

Figure 2. Forwarding probability curves in a random 6-regular network of 100,000 nodes

In order to compare with the other known strategies, our simulator implements, apart from our proposal, plain flooding, modified BFS [5] and ARPS [11]. Due to space limitation, we present experiments only with random and random regular p2p networks of 100.000 nodes. The average node degree varies from 4 to 6 and the resource popularity ranges as $0.0001 \leq q \leq 0.0003$, that is resources have from 10 to 30 replicas for the particular network size we consider, i.e. not very popular, stressing the search procedure. Table I summarizes the simulation parameters.

Before we present the comparative results, in Fig. 1–2 we present an example of the effectiveness of our approach. For this plot, we have considered a random 6-regular network.

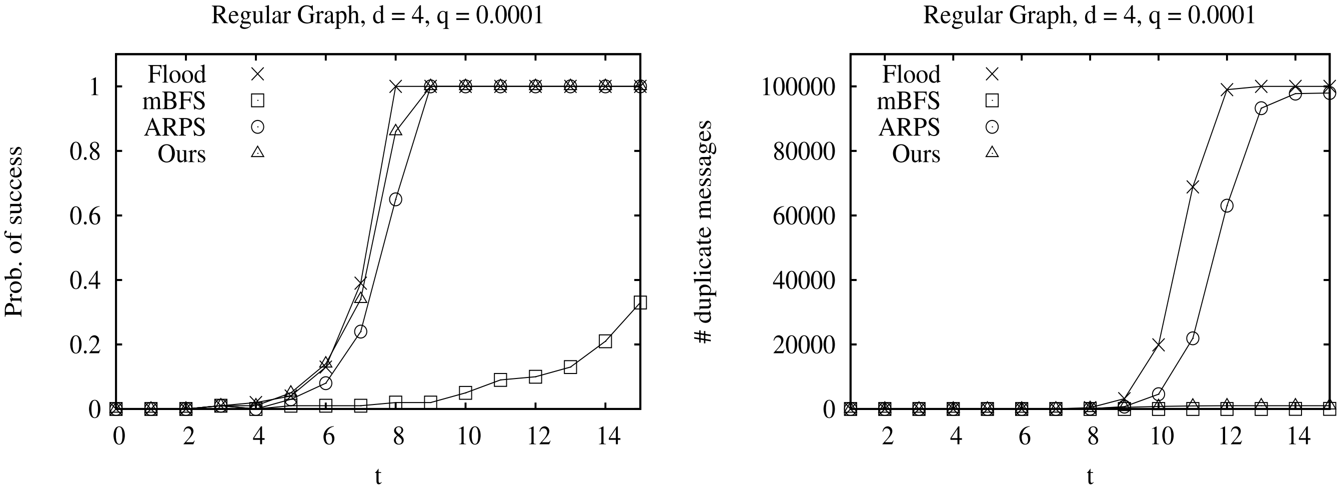

(a) (b)

Figure 3. Probability of success (a) and duplicate messages in random 4-regular network of 100,000 nodes and $q = 10^{-4}$.

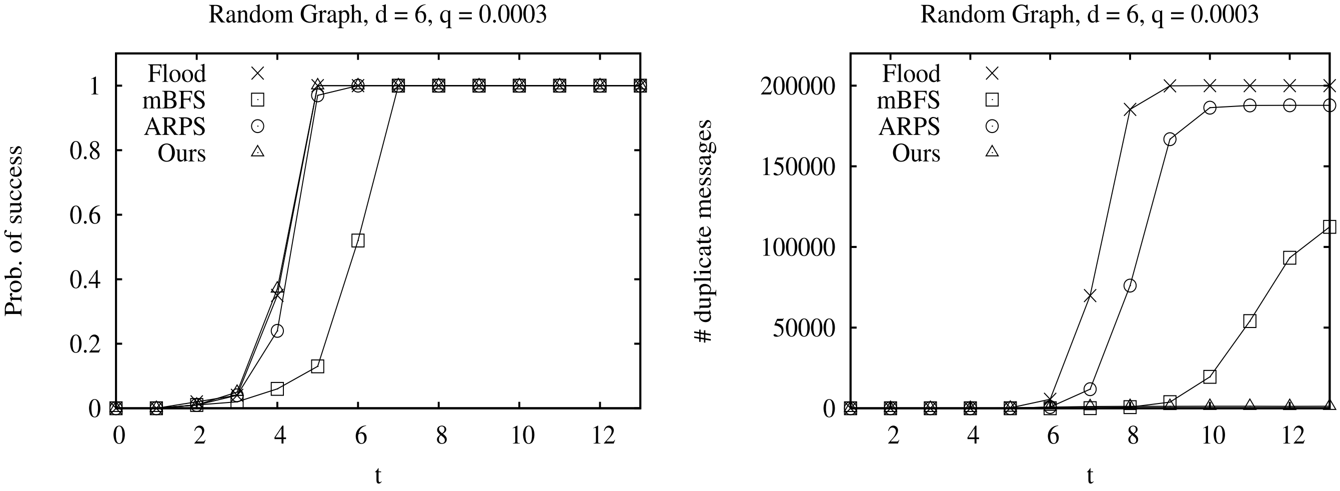

(a) (b)

Figure 4. Probability of success (a) and duplicate messages in random 6-regular network of 100,000 nodes and $q = 0.0003$.

In Fig. 1 we issue queries for a resource with popularity $q = 10^{-4}$. The curves show the number of covered nodes, as obtained from the simulation, and as predicted by (3), showing the accuracy of the estimation. In Fig. 2 we plot the forwarding probability for different values of resource popularity. Notice the fact that the shape of the curves are the exact opposite of the coverage curve in Fig. 1. The forwarding probability is initially high and then becomes low; the transition occurs at approximately the same step where coverage is becoming high. We believe this plays an important role for the success of our strategy.

Figs. 3–4 present the performance of the proposed strategy in comparison with the others. In Fig. 3 we have considered a 4-regular graph, where we search for a rare resource with popularity $q = 10^{-4}$, while in Fig. 4 the resource is more popular ($q = 3 \times 10^{-4}$) and the network is random, with an average node degree of 6. In all cases we give both the probability of success and the overhead, as the number of duplicate message transmissions.

In Fig. 3(a) we observe that most strategies manage to achieve 100% success probability with a TTL value of 9 hops, where most of the nodes have been visited. Notice however that our method converges faster and reaches close to 90% success rate already within 8 hops. At the same time, Fig. 3(b) shows that this is achieved with minimal duplicate message overheads. Modified BFS also has a small duplicate message cost but this is due to the low success rate it achieves as seen in Fig. 3(a).

The situation is similar in Fig. 4 where all strategies have better performance due to the higher popularity of the required resource. Even in this case, however, the superiority of our scheme is apparent since it combines the effective success rate with almost no duplicate message overheads, in contrast to all the other methods.

It should be clear that the performance of our strategy is quite impressive, outperforming the others both in term of success rate and in terms of overhead. It manages to approach the performance of pure flooding but without incurring its cost.

V. CONCLUSION

In this work, we present a novel probabilistic flooding strategy that can be employed for query routing in unstructured p2p networks. In contrast to other strategies, our approach is based on an estimate of the nodes that are covered at each step, and the popularity of the requested resource. We then utilize this estimate to derive the forwarding probability at each node that receives the query message and is about to transmit it further. We validate our estimates and evaluate experimentally our algorithm through detailed simulations. Moreover, we compare our strategy with other known approaches, showing that it offers high success probabilities combined with very low message overheads.

Our future plans include enhancing our strategy and expanding its applicability to other topologies such as power-law overlays. Finally, we are currently studying the effectiveness of our method in high-churn situations, with frequent node arrivals and departures.

REFERENCES

[1] V. V. Dimakopoulos and E. Pitoura, "On the performance of flooding-based resource discovery," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 11, pp. 1242–1252, 2006.

[2] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker, "Search and replication in unstructured peer-to-peer networks," in *ICS*, 2002, pp. 84–95.

[3] D. Tsoumakos and N. Roussopoulos, "Adaptive probabilistic search for peer-to-peer networks," in *Proceedings of the 3rd International Conference on Peer-to-Peer Computing*, ser. P2P '03. Washington, DC, USA: IEEE Computer Society, 2003, pp. 102–.

[4] C. Gkantsidis, M. Mihail, and A. Saberi, "Hybrid search schemes for unstructured peer-to-peer networks," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 3, march 2005, pp. 1526 – 1537 vol. 3.

[5] V. Kalogeraki, D. Gunopulos, and D. Zeinalipour-Yazti, "A local search mechanism for peer-to-peer networks," in *Proceedings of the eleventh international conference on Information and knowledge management*, ser. CIKM '02. New York, NY, USA: ACM, 2002, pp. 300–307.

[6] N. Makino, S. Arakawa, and M. Murata, "A flooding method for exchanging routing information in power-law networks," in *APCC 2005, Asia-Pacific Conference on Communications*, oct. 2005, pp. 812 –816.

[7] F. Banaei-Kashani and C. Shahabi, "Criticality-based analysis and design of unstructured peer-to-peer networks as "complex systems"," in *Proceedings of the 3st International Symposium on Cluster Computing and the Grid*, ser. CCGRID '03. Washington, DC, USA: IEEE Computer Society, 2003, pp. 351–.

[8] S. Crisóstomo, U. Schilcher, C. Bettstetter, and J. a. Barros, "Analysis of probabilistic flooding: how do we choose the right coin?" in *Proceedings of the 2009 IEEE international conference on Communications*, ser. ICC'09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 2080–2085.

[9] K. Oikonomou, D. Kogias, and I. Stavrakakis, "Probabilistic flooding for efficient information dissemination in random graph topologies," *Comput. Netw.*, vol. 54, pp. 1615–1629, July 2010.

[10] A. Stauffer and V. Barbosa, "Probabilistic heuristics for disseminating information in networks," *IEEE/ACM Transactions on Networking*, vol. 15, no. 2, pp. 425 –435, april 2007.

[11] H. Zhang, L. Zhang, X. Shan, and V. Li, "Probabilistic search in p2p networks with high node degree variation," in *ICC '07, IEEE International Conference on Communications, 2007*, june 2007, pp. 1710 –1715.

[12] R. Gaeta, G. Balbo, S. C. Bruell, M. Gribaudo, and M. Sereno, "A simple analytical framework to analyze search strategies in large-scale peer-to-peer networks," *Perform. Eval.*, vol. 62, no. 1-4, pp. 1–16, 2005.

[13] A. Klemm, C. Lindemann, M. K. Vernon, and O. P. Waldhorst, "Characterizing the query behavior in peer-to-peer file sharing systems," in *Proc. IMC '04, 4th ACM SIGCOMM conference on Internet measurement*, Sicily, Italy, 2004, pp. 55–67.