# Proactive gossip-based management of semantic overlay networks

Spyros Voulgaris[1, *, †], Maarten van Steen[2] and Konrad Iwanicki[2]

[1]*Department of Computer Science, ETH Zurich, Haldeneggsteig 4, 8092 Zurich,
Switzerland*
[2]*Department of Computer Science, Vrije Universiteit Amsterdam,
De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands*

SUMMARY

Much research on content-based P2P searching for file-sharing applications has focused on exploiting semantic relations between peers to facilitate searching. Current methods suggest *reactive* ways to manage semantic relations: they rely on the usage of the underlying search mechanism, and infer semantic relationships based on the queries placed and the corresponding replies received. In this paper we follow a different approach, proposing a *proactive* method to build a semantic overlay. Our method is based on an epidemic protocol that clusters peers with similar content. Peer clustering is done in a completely implicit way, that is, without requiring the user to specify preferences or to characterize the content of files being shared. In our approach, each node maintains a small list of semantically optimal peers. Our simulation studies show that such a list is highly effective when searching files. The construction of this list through gossiping is efficient and robust, even in the presence of changes in the network. Copyright © 2007 John Wiley & Sons, Ltd.

1. INTRODUCTION

File-sharing peer-to-peer (P2P) systems have gained enormous popularity in recent years. This has stimulated significant research activity in the area of content-based searching. Sparkled by the legal adventures of Napster, and challenged to defeat the inherent limitations concerning the scalability and

---

*Correspondence to: Spyros Voulgaris, Department of Computer Science, ETH Zurich, Haldeneggsteig 4, 8092 Zurich, Switzerland.
†E-mail: spyros@inf.ethz.ch

WILEY InterScience®
DISCOVER SOMETHING GREAT

failure resilience of centralized systems, research has focused on *decentralized* solutions for content-based searching, which by now has resulted in a wealth of proposals for peer-to-peer networks.

In this paper, we are interested in those groups of networks in which searching is based on grouping semantically related nodes. In these networks, a node first queries its semantically close peers before resorting to search methods that span the entire network. In particular, we are interested in solutions where semantic relationships between nodes are captured implicitly. This capturing is generally achieved through analysis of query results, leading to the construction of a local *semantic list* at each peer, consisting of references to other, semantically close peers.

Only very recently, an extensive study has been published on search methods in peer-to-peer networks, be they structured, unstructured, or of a hybrid form [1]. This study reveals that virtually all peer-to-peer search methods in semantic overlay networks follow an integrated approach towards the construction of the semantic lists, while at the same time accounting for changes occurring in the set of nodes. These changes involve the joining and leaving of nodes, as well as changes in a node's preferences.

The problem we are faced with is that the construction of semantic lists should result in highly clustered overlay networks. These networks excel for searching content when nothing changes. However, handling dynamics requires the discovery and propagation of changes that may happen *anywhere* in the network. For this reason, overlay networks should also reflect desirable properties of random graphs and complex networks in general [2,3]. These two conflicting demands generally lead to complexity when integrating solutions into a single protocol.

Protocols for content-based searching in peer-to-peer networks should separate these concerns. In particular, we advocate that when it comes to constructing and using semantic lists, these lists should be optimized for search only, regardless of any other desirable property of the resulting overlay. Instead, a separate protocol should be used to handle network dynamics, and provide up-to-date information that will allow proper adjustments in the semantic lists (and thus leading to adjustments in the semantic overlay network itself).

In this paper we propose such a two-layered approach for managing semantic overlay networks. The top layer contains a gossip-based protocol that strives to optimize semantic lists for searching only. The bottom layer offers a fully decentralized service for delivering, in an unbiased fashion, information on new events, similar in nature to the peer-sampling service described in [4]. Our main contribution is that we demonstrate that this two-layered approach leads to high-quality semantic overlay networks. We substantiate our claims through extensive simulations using traces collected from the eDonkey file-sharing network [5].

The paper is organized as follows. We start with presenting our protocols in the next section, followed by describing our experimental setup in Section 3. Performance evaluation is discussed in Section 4, followed by an analysis of consumed bandwidth in Section 5. We conclude with a discussion in Section 6.

2. THE PROTOCOL

In our model each peer maintains a dynamic list of semantic neighbors, called its *semantic view*, of fixed small size ℓ. A peer searches for a file by first querying its semantic neighbors. If no results are returned, the peer then resorts to the default search mechanism.

2.1. Outline

Our aim is to organize the semantic views so as to maximize the hit ratio of the first phase of the search. We will call this the *semantic hit ratio*. We anticipate that the probability of a neighbor satisfying a peer's query is proportional to the semantic proximity between the peer and its neighbor. We aim, therefore, at filling a peer's semantic view with its ℓ semantically closest peers out of the whole network.

We assume the existence of a *semantic proximity function* $S(F_P, F_Q)$, which given the file lists F_P and F_Q of peers P and Q, respectively, provides a numeric metric of the semantic proximity between the two peers. The more semantically similar the file lists of P and Q are, the higher the value of $S(F_P, F_Q)$. We are essentially seeking to pick peers $Q_1, Q_2, \ldots, Q_\ell$ for peer P's semantic view, such that the sum $\sum_{i=1}^{\ell} S(F_P, F_{Q_i})$ is maximized.

We assume that the semantic proximity function exhibits some sort of transitivity, in the sense that if P and Q are semantically similar to each other, and so are Q and R, then some similarity between P and R is likely to hold. Note that this transitivity does not constitute a hard requirement for our system. In its absence, semantically related neighbors are discovered based on random encounters. If it exists though, it is exploited to dramatically enhance efficiency.

2.2. Design motivation

From our previous discussion, we are seeking means to construct, for each node, a semantic view from all the current nodes in the system. There are two sides to this construction.

First, based on the assumption of transitivity in the semantic proximity function S, a peer should explore the semantically close peers that its neighbors have found. In other words, if Q is in P's semantic view, and R is in Q's view, it makes sense to check whether R is also semantically close to P. Exploiting the transitivity in S should then quickly lead to high-quality semantic views.

Second, it is important that *all* nodes are examined. The problem with following only transitivity is that we eventually will be searching only in a single semantic cluster. Similar to the special 'long' links in small-world networks [6], we need to establish links to *other* semantically related clusters. Likewise, when new nodes join the network, they should easily find an appropriate cluster to join. These issues call for a randomization when selecting nodes to inspect for adding to a semantic view.

In our design we decouple these two aspects by adopting a two-layered set of gossip protocols, as can be seen in Figure 1. The lower layer, called CYCLON [7], is responsible for maintaining a connected overlay and for periodically feeding the top-layer protocol with nodes uniformly randomly selected from the network. In its turn, the top-layer protocol, called VICINITY, is in charge of focusing on discovering peers that are semantically as close as possible, and of adding these nodes to the semantic views.

2.3. Gossiping framework

All information exchanges between peers are carried out by means of *gossip items*, or simply *items*. A gossip item created by peer P is a tuple containing the following three fields:

1. P's contact information (network address and port).
2. The item's creation time.
3. Application-specific data; in this case P's file list.

Figure 1. The two-layered framework.

```
/*** Active thread ***/
// Runs periodically every T time units
q = selectPeer()
myItem = (myAddress, timeNow, myFileList)
buf_send = selectItemsToSend()
send buf_send to q
receive buf_recv from q
view = selectItemsToKeep()

/*** Passive thread ***/
// Runs when contacted by some peer
receive buf_recv from p
myItem = (myAddress, timeNow, myFileList)
buf_send = selectItemsToSend()
send buf_send to p
view = selectItemsToKeep()
```

Figure 2. Epidemic protocol skeleton.

Each node maintains locally a number of items per protocol, called the protocol's *view*. This number is the same for all items, and is called the protocol's *view size* (c_v for VICINITY, and c_c for CYCLON).

Figure 2 presents a generic skeleton forming the basis for both VICINITY and CYCLON gossiping protocols. Each node runs two threads. An *active* one, which periodically wakes up and initiates communication to another peer, and a *passive* one, which responds to the communication initiated by another peer.

The functions appearing underlined, namely selectPeer(), selectItemsToSend(), and selectItemsToKeep() form the three *hooks* of this skeleton. Different protocols can be instantiated from this skeleton by implementing specific policies for these three functions, in turn, leading to different emergent behaviors.

Hook	Description
selectPeer()	Select item with the oldest timestamp
selectItemsToSend()	
RANDOM	Randomly select g_v items
SELECTIVE	Select the g_v items of nodes semantically closest to the selected peer
COMPLETE	Select the g_v items of nodes semantically closest to the selected peer from the VICINITY view *and* the CYCLON view
selectItemsToKeep()	Keep the c_v items of nodes that are semantically *closest*, out of items in its current view, items received, and items in the local CYCLON view. In case of multiple items from the same node, keep the one with the most recent timestamp.

(a)

Hook	Description
selectPeer()	Select item with the oldest timestamp
selectItemsToSend():	
active thread	Select own item and randomly $g_c - 1$ others from the CYCLON view
passive thread	Randomly select g_c items from the CYCLON view
selectItemsToKeep():	
active thread	Keep all g_c received items, replacing (if needed) the item selected by selectPeer() and the $g_c - 1$ ones selected to send.
passive thread	Keep all g_c received items, replacing (if needed) the g_c ones selected to send.
	In case of multiple items from the same node, keep the one with the most recent timestamp.

(b)

Figure 3. The chosen policies for: (a) the VICINITY protocol and (b) the CYCLON protocol.

The number of items exchanged in each communication is predefined, and is called the protocol's *gossip length* (g_v for VICINITY, and g_c for CYCLON).

For VICINITY, we chose the policies shown in Figure 3(a). Note that selectItemsToKeep() takes into account CYCLON's cache too in selecting the best c_v items to keep. This is the default link between the two layers. As we discuss below, COMPLETE will turn out to be an excellent choice for forming semantic clusters.

For CYCLON, we made the choices shown in Figure 3(b). CYCLON is a protocol we previously developed, and which is extensively described and analyzed in [7]. Effectively, what selectItemsToSend() and selectItemsToKeep() establish is an *exchange* of some neighbors between the caches of the two communicating peers. In addition to that, the selected peer's item in the initiator's cache is always removed, but the initiator's (new) item is always placed in the selected peer's cache.

CYCLON creates an overlay with completely random, uncorrelated links between nodes, such that the in-degree (the number of incoming links) is practically the same for each node. Importantly, it can achieve this property fairly quickly even when a small number of items (such as 3 or 4) is exchanged in each communication, even for large caches of several dozens of items. Therefore, it

is ideal as a lightweight service that can offer a node a randomly selected peer from the current set of nodes.

3. EXPERIMENTAL ENVIRONMENT AND SETTINGS

All experiments presented here have been carried out with PeerSim [8], an open source simulator for P2P protocols, developed in Java at the University of Bologna.

To evaluate our protocol, we used real-world traces from the eDonkey file-sharing system [9], collected by Le Fessant *et al.* in November 2003 [5]. A set of 11 872 worldwide distributed peers along with the files each one shares is logged in these traces. A total number of 923 000 unique files is being collectively shared by these peers.

In order to simplify the analysis of our system's emergent behavior, we determined equal gossiping periods for both layers. More specifically, once every T time units each node initiates first a gossip exchange with respect to its bottom (CYCLON) layer, immediately followed by a gossip exchange at its top (VICINITY) layer. Note that even though nodes initiate gossiping at universally fixed intervals, they are not synchronized with each other.

Even though both protocols are asynchronous, it is convenient to introduce the notion of *cycles* in order to study their evolutionary behavior with respect to time. We define a cycle to be the time period during which *each* node has initiated gossiping exactly *once*. Since each node initiates gossiping periodically, once every T time units, a cycle is equal to T time units.

A number of parameters had to be set for these experiments, listed here.

Proximity function S: We chose a rather simple, yet intuitive proximity function to test our protocol with. The proximity S between two nodes P and Q, with file lists F_P and F_Q, respectively, is defined as the number of files that lay in both lists. More formally: $S(F_P, F_Q) = |F_P \cap F_Q|$. As stated in 2.1, semantically closer the two nodes are, the higher the value of S is. Note that our goal was to demonstrate the power of our gossiping protocol in forming a semantic network based on *a* proximity function. Even though much richer proximity functions could have been applied, it was out of the scope of this paper.

Semantic view size ℓ: In all experiments the semantic view consisted of the 10 semantically closest peers in the VICINITY cache. As shown in [10], a semantic view size of $\ell = 10$ provides a good tradeoff between the number of nodes contacted in the semantic search phase and the expected semantic hit ratio.

Cache size: For the cache size selection, we are faced with the following tradeoff for both protocols. A large cache size provides higher chances of making better item selections, and therefore accelerate the construction of (near-)optimal semantic views. On the other hand, the larger the cache size, the longer it takes to contact all peers in it, resulting in the existence of older—and therefore more likely to be invalid—links. Of course, a larger cache also takes up more memory, although this is generally not a significant constraint nowadays.

Considering this tradeoff, and based on experiments not further described here, we fixed the cache size to be 100 as a basis to compare different configurations. When both VICINITY and CYCLON are used, they are allocated 50 cache entries each.

Gossip length: The gossip length, that is, the number of items gossiped per gossip exchange per protocol, is a crucial factor for the amount of bandwidth used. This becomes of greater consequence,

considering that an item carries the file list of its respective node. So, even though exchanging more items per gossip exchange allows information to disseminate faster, we are inclined to keep the gossip lengths as low as possible, as long as the system's performance is reasonable.

Again, for the sake of comparison, we fixed the total gossip length to six items. When both VICINITY and CYCLON are used, each one is assigned a gossip length of 3.

Gossip period T: The gossip period is a parameter that does not affect the protocol's behavior. The protocol evolves as a function of the number of messages exchanged, or, consequently, of the number of cycles elapsed. The gossip period only affects how fast the protocol's evolution will take place in time. The single constraint is that the gossip period T should be adequately longer than the worse latency throughout the network, so that gossip exchanges are not favored or hindered due to latency heterogeneity. A typical gossip period for our protocol would be 1 min, even though this does not affect the following analysis.

4. PERFORMANCE EVALUATION

4.1. Convergence speed on cold start

To evaluate the convergence speed of our algorithm, we first test how quickly it groups semantically related peers, when starting with a semantically unaware network.

The objective, as imposed by the proximity function, is for each node to discover the ℓ peers that have the most common files with it. We define a node's *semantic view quality* to be the ratio of the number of common files shared with its current ℓ semantic neighbors over the number of common files it would share with its ℓ optimal semantic neighbors.

Figure 4(a) shows the average semantic view quality as a function of the cycle for four distinct configurations. In favor of comparison fairness, the cache size and gossip length are 50 and 3, respectively, in each layer, for all configurations. The only exception is the first configuration, which
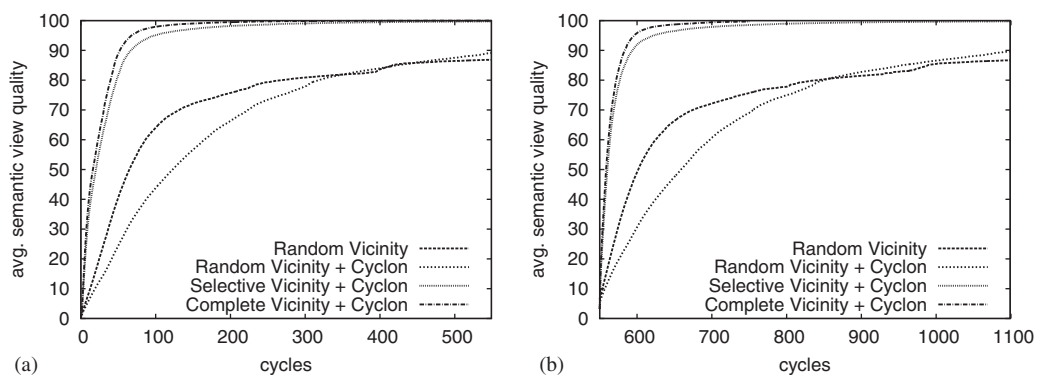

Figure 4. (a) Convergence of semantic views' quality and (b) evolution of semantic views' quality for a sudden change in all users' interests at cycle 550.

has a single layer. In this case, the cache size and gossip length are 100 and 6, respectively. All experiments start with each node knowing 5 random other ones, simply to ensure initial connectivity in a single connected cluster.

In the first configuration, RANDOM VICINITY is running stand-alone. The progress of the semantic views' quality is rather steep in the first 100 cycles, but as nodes gradually concentrate on their very own neighborhood, getting to know new, possibly better peers becomes rare, and progress slows down.

In the second configuration, a two-layered approach consisting of RANDOM VICINITY and CYCLON is running. The slow start compared to stand-alone VICINITY is a reflection of the smaller VICINITY cache (3 as opposed to 6). However, the two-layered approach's advantage becomes apparent later, when CYCLON keeps feeding the RANDOM VICINITY layer with new, uniformly randomly selected nodes, maintaining a higher progress rate, and outperforming stand-alone VICINITY in the long run.

In the third configuration, SELECTIVE VICINITY demonstrates its contribution, as progress is significantly faster in the initial phase of the experiment. This is to be expected, since the items sent over in each SELECTIVE VICINITY communication, are the ones that have been selected as the semantically closest to the recipient.

Finally, in the fourth configuration, COMPLETE VICINITY keeps the progress rate high even when the semantic views are very close to their optimal state. This is due to the broad random sampling achieved by this version. In every communication, a node is exposed to the best peers out of 50 *random* ones, in addition to 50 peers from its neighbor. In this way, semantically related peers that belong to separate semantic clusters quickly discover each other, and subsequently the two clans merge into a single cluster in practically no time.

4.2. Adaptivity to changes of user interests

In order to test our protocol's adaptivity to dynamic user interests, we ran experiments where the interests of some users changed. We simulated the interest change by picking a random pair of nodes and swapping their file lists in the middle of the experiment. At that point, these two nodes found themselves with semantic views unrelated to their (new) file lists, and therefore had to gradually climb their way up to their new semantic vicinity, and replace their useless links by new, useful ones.

Once again, we present the worst case—practically unrealistic—scenario, of *all* nodes changing interests at once, at cycle 550 of the experiment of Figure 4(a). The evolution of the average semantic view quality from the moment when all nodes change interests is presented in figure 4(b). The faster convergence compared to Figure 4(a) is due to the fact that views are already fully filled up at cycle 550, so nodes have more choices to start looking for good candidate neighbors.

Even though this scenario is very unrealistic, it demonstrates the power of our protocol in adapting to even massive scale changes. This adaptiveness is due to the priority given to newer items in `selectItemsToKeep()`, which allows a node's items with updated semantic information to replace older items of that node fast.

4.3. Effect on semantic hit ratio

In order to further substantiate our claim that semantic-based clustering endorses P2P searching, we conducted the following experiments. A randomly selected file was removed from *each* node,

Figure 5. Semantic hit ratio, for gossip lengths 1, 3, and 5 in each layer.

and the system was run considering proximity based on the remaining files. Then, each node did a search on the file it was missing. We measured the semantic hit ratio to be over 36% for a semantic view of size 10.

Figure 5 presents the semantic hit ratio as a function of the cycle. Three experiments are shown, with gossip lengths for *both* layers set to 1, 3, and 5. Note that computation of the hit ratio for each cycle was made offline, without affecting the mainstream experiment's state.

4.4. Behavior under changing membership

To investigate the behavior of our algorithm as nodes join and leave the network, we conducted two different experiments with COMPLETE VICINITY. First, we are interested to see how fast a node discovers optimal neighbors for its semantic view, when joining an already converged network.

To this end, we conduct a series of experiments, each time starting with a network from which one randomly selected node is removed. After a network has converged, we add the missing node and measure the number of cycles it takes to fill, respectively 50, 90, and 100% of its semantic view with *optimal* neighbors. The respective cumulative distribution function (CDF) graphs are shown in Figure 6.

The experiments clearly show that the semantic view is rapidly filled with optimal neighbors for the vast majority of nodes, although some may take considerably more time. It can also be seen that, although discovering *all* best neighbors may take arguably long for some nodes, it takes significantly fewer cycles to discover *most* best neighbors (CDFs for finding 50 and 90% of best neighbors shown).

Let us now take a look at how the algorithm behaves under churn, that is, when nodes regularly join and leave the network. For this experiment, we consider an initial converged network of 10 000 nodes. During each cycle we remove n nodes and replace them with n other nodes. Every node in the system corresponds to one node from the eDonkey traces (i.e. stores the same files), which contained a total of 11 872 nodes. Therefore, at any moment in time a random subset of 10 000

out of 11 872 nodes is active, and the remaining 1872 nodes are down. The analysis of churn rates from real-world Gnutella traces [11] shows that the set of active nodes changes by approximately 0.2% every 10 s. In other words, if we assume a cycle length of 10 s, a realistic value for n is 20. We have also experimented with larger churn rates, namely 1 and 5%, which correspond to a cycle duration of 50 and 250 s, respectively.

The results of these experiments are shown in Figure 7(a) in which the average number of optimal alive neighbors in semantic views under different churn values is plotted. We see that as the churn rate increases, the semantic views generally remain polluted with links to non-optimal neighbors. This can be easily explained by considering the VICINITY cache (from which the neighbors for the


Figure 6. CDF of the speed by which the semantic view of a joining node is filled with optimal neighbors.
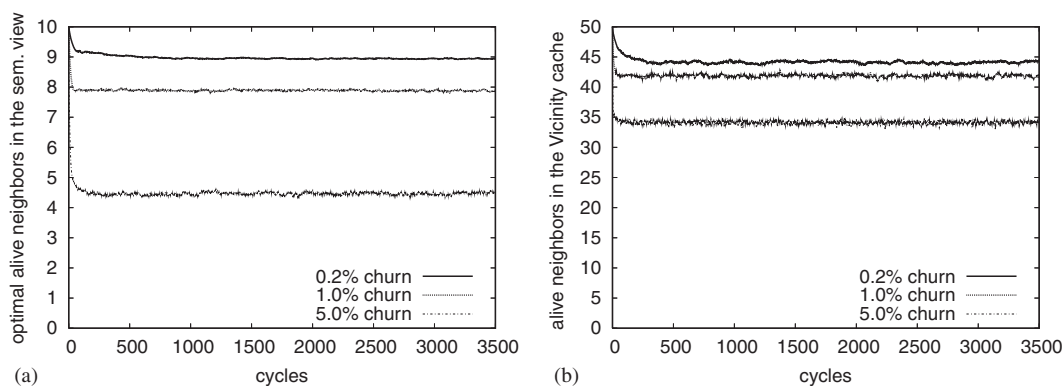

Figure 7. (a) The average number of optimal alive neighbors in the semantic view and (b) the average number of alive neighbors in the VICINITY cache.

semantic views are extracted). In Figure 7(b), we see that under high churn the cache contains a relatively large fraction of links to dead nodes. As these inactive links may refer to nodes with high number of common files, they prevent establishing optimal links with nodes that are alive but share a smaller number files. Moreover, when a node is reborn, it can take some time for other nodes, which now may have non-optimal semantic views, to establish links to it (cf. Figure 6).

5. BANDWIDTH CONSIDERATIONS

Due to the periodic behavior of gossiping, the price of having rapidly converging and accurate protocols may inhibit a high usage of network resources (i.e. bandwidth).

In each cycle, a node gossips on average twice (exactly once as an initiator, and on average once as a responder). In each gossip $2 \cdot (g_v + g_c)$ items are transferred to and from the node, resulting in a total traffic of $4 \cdot (g_v + g_c)$ items for a node per cycle. An item's size is dominated by the file list it carries. A single file is identified by its 128-bit (16-byte) MD4 hash value. Analysis of the eDonkey traces [5] revealed an average number of 100 files per node (more accurately, 99.35). Therefore, a node's file list takes on average 1600 bytes. So, in each cycle, the total number of bytes transferred *to* and *from* the node is $6400 \cdot (g_v + g_c)$.

For $g_v = g_c = 3$, the average amount of data transferred to and from a node in one cycle is 38 400 bytes, while for $g_v = g_c = 1$, it is just 12 800. Maintaining almost optimal semantic views requires frequent gossiping to account for the churn. Based on the traces, to achieve 90% optimality of the semantic view, the churn rate must be limited to approximately 0.2%. Consequently, the gossip period T must be equal to 10 s, which translates to an average bandwidth of 3840 bytes per second for $g_v = g_c = 3$, and 1280 bytes per second for $g_v = g_c = 1$. However, if 80% optimality is acceptable, the gossip period can be reduced to 50 s, which yields 768 and 256 bytes per second, respectively, a factor of 5 improvement traded for only 10% quality degradation.

We consider such a bandwidth consumption to be rather small, if not negligible compared with the bandwidth used for the actual file downloads. It is, in fact, a small price to pay for relieving the default search mechanism from about 35% of the search load, which is often significantly higher (e.g. flooding or random-walk search). Moreover, the bandwidth can be further reduced by employing techniques such as Bloom filters [12] and on-demand fetching of file lists, instead of associating a full file list with each cache item.

6. DISCUSSION

To the best of our knowledge, all earlier work on implicit building of semantic overlays relies on using heuristics to decide *which* of the peers that served a node recently are likely to be useful again in future queries [10,13,14].

However, all these techniques inhibit a weakness that challenges their applicability to the real world. They all assume a *static* network, free of node departures, which is a rather strong assumption considering the highly dynamic nature of file-sharing communities. Also, it is not clear how they perform in the presence of dynamic user preferences.

Moreover, as opposed to the existing solutions, our algorithm can, to some extent, help against so-called free-riders in the P2P file-sharing networks [15]. Free-riders provide no files to be downloaded by other users, but still use the network to obtain files that are interesting for themselves. Because of a lack of shared files, VICINITY does not create any meaningful semantic views for such misbehaving nodes. This, combined with heuristics forbidding frequent usage of the backup search algorithms, minimizes the number of successful searches for free-riders, and consequently discourages the free-riding practice.

Regarding proximity-based P2P clustering, our work comes close to the T-Man protocol [16], which has been developed independently. Although there were significant differences with the original T-Man protocol, the most recent version shows a strong similarity with our work. An important difference remains that VICINITY and CYCLON offer the means to control bandwidth by arbitrarily deciding on the number of entries that need to be exchanged. More important, however, is that we show in this paper how our specific two-layered gossiping approach can be successfully applied to searching in peer-to-peer file-sharing systems. Such an evaluation has not yet been done before.

Concluding, in this paper we introduced the idea of applying epidemics to proactively build and dynamically maintain semantic lists in a large-scale file-sharing system. Specifically, we showed that using a two-layered approach combining two epidemic protocols is the appropriate way to build such a service. Finally, we presented a fast converging, highly adaptable, yet lightweight epidemic-style solution to this problem.

ACKNOWLEDGEMENTS

REFERENCES

1. Risson J, Moors T. Survey of research towards robust peer-to-peer networks: Search methods. *Technical Report UNSW-EE-P2P-1-1*, University of New South Wales, Sydney, Australia, September 2004.
2. Albert R, Barabasi A-L. Statistical mechanics of complex networks. *Reviews of Modern Physics* 2001; **74**(1):47–97.
3. Newman MEJ. Random graphs as models of networks. *Handbook of Graphs and Networks*: *From the Genome to the Internet*, ch. 2, Bornholdt S, Schuster HG (eds.). Wiley: New York, NY, 2002.
4. Jelasity M, Guerraoui R, Kermarrec A-M, van Steen M. The peer sampling service: Experimental evaluation of unstructured gossip-based implementations. *Middleware 2004* (*Lecture Notes in Computer Sciences*, vol. 3231). ACM/IFIP/USENIX. Springer: Berlin, 2004.
5. Fessant F, Handurukande S, Kermarrec A-M, Massoulié L. Clustering in peer-to-peer file sharing workloads. *Third International Workshop on Peer-to-Peer Systems* (*IPTPS*), San Diego, U.S.A., February 2004.
6. Watts DJ. *Small Worlds. The Dynamics of Networks between Order and Randomness*. Princeton University Press: Princeton, NJ, 1999.
7. Voulgaris S, Gavidia D, van Steen M. Cyclon: Inexpensive membership management for unstructured p2p overlays. *Journal of Network and Systems Management* 2005; **13**(2):197–217.
8. PeerSim. http://peersim.sourceforge.net/ [July 2007].
9. eDonkey. http://www.edonkey2000.com/ [November 2003].
10. Handurukande S, Kermarrec A-M, Le Fessant F, Massoulié L. Exploiting semantic clustering in the edonkey p2p network. *Eleventh ACM SIGOPS European Workshop* (*SIGOPS*), Leuven, Belgium, September 2004.
11. Saroiu S, Gummadi KP, Gribble SD. Measuring and analyzing the characteristics of Napster and Gnutella hosts. *ACM Multimedia Systems* 2003; **9**(2):170–184.

12. Bloom BH. Space/time trade-offs in hash coding with allowable errors. *Communications of ACM* 1970; **13**(7):422–426.
13. Sripanidkulchai K, Maggs B, Zhang H. Efficient content location using interest-based locality in peer-to-peer systems. *INFOCOM Conference*, San Francisco, U.S.A., 2003.
14. Voulgaris S, Kermarrec A, Massoulié L, van Steen M. Exploiting semantic proximity in peer-to-peer content searching. *Tenth International Workshop on Future Trends in Distributed Computing Systems* (*FTDCS 2004*), Suzhu, China, November 2001.
15. Adar E, Huberman BA. Free riding on gnutella. *First Monday* 2000. Avaiable at: http://firstmonday.org/issues/issue5_10/adar/.
16. Jelasity M, Babaoglu O. T-Man: Gossip-based overlay topology management. In *Engineering Self-Organising Systems: Third International Workshop (ESOA 2005) (Lecture Notes in Computer Science*, vol. 3910*)*, Brueckner SA, Di Marzo Serugendo G, Hales D, Zambonelli F (eds.), Revised Selected Papers. Springer-Verlag, 2006; 1–15. DOI: 10.1007/11734697_1.