

A Two-Level Semantic Caching Scheme for Super-Peer Networks

Paweł Garbacki and Dick H. J. Epema
Faculty of Electrical Engineering,
Mathematics, and Computer Science
Delft University of Technology,
P.O. Box 5031, 2600 GA Delft,
The Netherlands
{p.garbacki,d.h.j.epema}@ewi.tudelft.nl

Maarten van Steen
Department of Computer Science
Vrije Universiteit Amsterdam,
De Boelelaan 1081a, 1081 HV Amsterdam,
The Netherlands
steen@cs.vu.nl

Abstract

Recent measurement studies of file-sharing peer-to-peer networks have demonstrated the presence of semantic proximity between peers and between shared files. In this paper we tackle the problem of exploiting the semantic locality of peer requests to improve the performance of a P2P network by the use of semantic caches. Such caches group together peers with similar interests as well as files with similar request patterns. The resulting two-level caching infrastructure can be integrated in a very natural way with the super-peer concept. The super-peers in our system cache pointers to files recently requested by their client peers. The client peers, on the other hand, keep caches of super-peers that answered most of their requests in the past. As a consequence, peers with similar interests are grouped (clustered) under the same super-peers, and also files with similar request patterns are indexed by the same super-peers. We discuss the design choices and optimizations of the presented model. We also evaluate our system versus the network that uses only one level of semantic caches.

1. Introduction

Peer-to-peer content sharing has become very popular in the last few years, and is nowadays the biggest consumer of Internet bandwidth [19, 23]. The success of P2P technology has attracted the interest of the research community, which has resulted in many improvements to the existing protocols as well as to completely new P2P designs [8, 15, 17]. However, the dominating P2P content sharing networks seem to ignore the dynamic properties of the shared data such as file popularities and flash crowds. Search is either blind [1, 3], independent of the query, or based on static aspects of the content [21, 22, 25, 29] such as file hashes. In this paper we

propose an optimization that can be integrated with both these searching schemes improving their performance by taking advantage of the semantic correlation [10, 12] between peers as well as shared files. The problem that needs to be solved here is the efficient management of semantic information. We tackle this issue by introducing a hierarchy into the network in the form of super-peers.

Super-peer networks [28] occupy the middle-ground between centralized and entirely symmetric P2P networks. Super-peers are selected nodes with extra capabilities, but also extra duties in the network. A super-peer acts as a server to a dynamic subset of weak (ordinary, client) peers. Weak peers submit queries to their super-peers and receive results from them. Super-peers are connected to each other forming an overlay network of their own, submitting and answering requests on behalf of the weak peers.

Creating connections between weak peers and super-peers can be seen as peer clustering; a cluster in this context is simply the set of peers that are connected to the same super-peer. The semantic structure in the shared content and the patterns observed in successive searches made by individual nodes can be used as the clustering criterion. Grouping together peers with similar interests under one super-peer increases query locality, and as a consequence improves the performance of search [18]. Clustering may be applied not only to the peers, but also to the content [9, 16, 24]. In a super-peer architecture with content clustering pointers to the files with similar request patterns are stored at the same super-peer.

In this paper we combine both peer- and content-based clustering. The proposed architecture assumes the existence of caches of two types. Weak peers keep lists of super-peers that proved in the past to be in some way most suitable for them. Super-peers index files residing at the weak peers that were recently requested by some peers.

The rest of the paper is organized as follows. Section 2 describes in detail the system architecture. Section 3 intro-

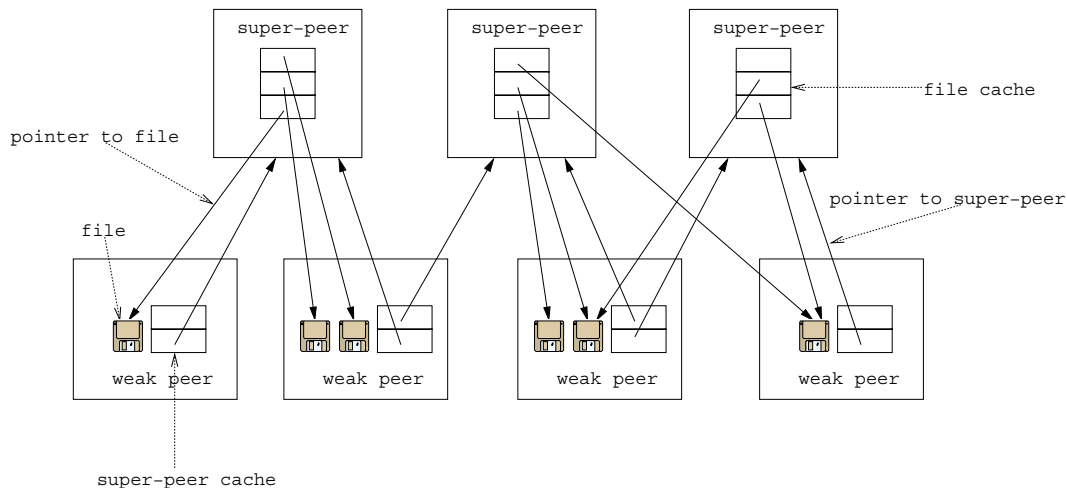


Figure 1. The two-level caching scheme.

duces a model of P2P network with semantic relationships between peers and files. This model is further used in Section 4 to evaluate the performance of our system. The paper concludes in Section 5 by exploring some opportunities for future work.

2. Detailed system model

In this section we describe in detail the concept of our P2P network based on super-peers and two-level semantic caches.

2.1. Design considerations

We take several design decisions solving the important problems hampering the performance of P2P architectures exploiting the semantic relationships between peers and content items. For each decision we provide a brief rationale for adopting it.

Caching is the most commonly used technique for discovering and profiting from the semantic dependencies in the P2P network. Most of the existing P2P infrastructures based on caching of the search results assume the existence of homogeneous caches at all nodes. Such approaches seem to ignore the fact that caching can be efficient only if a sufficient number of queries is generated. Each cache has a warm-up phase in which it collects information about the access characteristics of the stored items. A new peer joining the system has to first build its own cache before it can profit from the semantic structure of the network. We solve this problem by placing shared caches at a set of selected (super-)nodes. These caches are used (shared) by many peers at the same time. A peer that joins the system is auto-

matically bound to one or more super-peers and can immediately use the information collected by these super-peers.

The question that arises then is how replacing a number of local caches with one shared cache will influence the hit ratios? According to [7], the average hit ratio of a cache capable of storing n items is then proportional to $\log(n)$. Let's assume that we replace caches of size n located at m peers with a single shared cache of size n . The worst-case scenario occurs when peer interests are distinct (peers search completely different files). Assuming that each peer contacts the shared cache with the same frequency, the cache hit ratio is proportional to $\log(n/m)$, which can be written as $\log(n) - \log(m)$, because each peer uses a portion of size n/m of the shared cache. We conclude that the cache hit ratio decreases logarithmically with m . It should be added that in a real system the size of the shared cache is higher than n . Furthermore, many peers share interest for the same files, so we can expect much higher hit ratios.

Introducing super-peers into the architecture of a P2P system usually implies a serious limitation of the weak peer autonomy. We try to decrease the extent of this problem by establishing loose relationships between the weak peers and the super-peers. The weak peers are free to decide which super-peers they are connected to. This decision is made locally and independently from other peers. The super-peer selection rule has however a global property of grouping peers with similar interests under one super-peer.

2.2. System architecture

In this section we describe the architecture of our super-peer network incorporating a two-level caching scheme.

Figure 1 presents the data structures used in our system. The information stored at a node depends on the type of

this node. Each weak peer p has a *super-peer cache* $p.S$, which contains the identities of super-peers (e.g., their IP addresses and port numbers). Each super-peer s has a *file cache* $s.F$ of pointers to files stored at some peers.

All items in the super-peer and file caches are assigned *priorities*, which are non-negative integer numbers. The priority determines the importance of a particular item, the higher the better. The initial priority assigned to a data item when it is added to the cache and the way the priority is modified upon a cache hit are determined by the caching policy (see Section 2.5). There are two situations when the priorities are taken into account. First, when the cache capacity is exceeded, the item with the lowest priority is removed. Second, the priorities are used for optimizing query routing. Details are presented in Section 2.3.

The super-peers are interconnected with one of the standard P2P networks, which we leave unspecified. We require however that the probability that the search succeeds is high if the requested information is possessed by only one of the super-peers. Examples of protocols satisfying this criterion are Gnutella and epidemic-based approaches such as SCAMP [11].

Whenever a weak peer initiates a search, it first checks the file caches of the super-peers known to it. If the file is not found in one of these caches, a system-wide search in the super-peer network is initiated. The pointer to the located file is then cached by one of the super-peers known to the weak peer that initiated the search. Note that in this scheme only the files that are known to at least one of the super-peers can be located. To increase the probability of finding the file we could perform a search in the entire P2P network instead of among the super-peers only. As we show in Section 4.2, such an expanded search is usually not needed because most of the files are indexed by at least one super-peer.

The performance of the search is determined by the cache hit ratio. Peer p should have in its cache super-peers indexing content which is closest to p 's interests. This is achieved by preferring super-peers that provided in the past the highest number of positive responses to the queries submitted by p .

By allowing super-peers to cache pointers to the files recently requested by their weak peers, we exploit a simple, yet powerful principle called *interest-based locality* [24]. Interest-based locality postulates that if two peers are interested in the same file, it is very likely that more of their requests will overlap.

2.3. The search protocol

The following pseudocode describes the actions taken by peer p searching for file f :

```

peer_search(p : peer, f : file_name)
1: for (s in p.S ordered according to
    decreasing priorities)
2:   q := super-peer_local_search(s,f)
3:   if (super-peer_local_search succeeded)
    // f was found
4:     t := s
5:     break
6:   if (f was not yet found)
7:     s := super-peer in p.S selected
    randomly with probability
    proportional to its priority in p.S
8:   <q,t> := super-peer_search(s,f)
9:   if (super-peer_search did not succeed)
    // f was not found
10:    return ERROR "File f not found"
11: if (p.S contains t) then
12:   increase the priority of t in p.S
13: else
14:   insert t into p.S
15: return q // peer that has file f

```

In the above algorithm, peer p first tries to locate file f in the local caches of the super-peers in its super-peer cache. Note that p starts from the super-nodes with the highest priorities. If the file was not found in this way (line 6), the search request is forwarded to one of the super-peers in $p.S$ selected randomly based on its priority (line 7). This super-peer takes care of locating file f and returns a pair $\langle q, t \rangle$, where q is a peer that shares f and t is a super-peer that has a pointer $\langle f, q \rangle$ in its cache. Finally, peer p updates the priority of t in its cache (lines 11–14).

The algorithm of the super-peer local search is straightforward:

```

super-peer_local_search(s : super-peer,
    f : file_name)
1: if (an entry <f,q> exists in cache s.F)
2:   increase the priority of <f,q> in s.F
3:   return q
4: else
5:   return ERROR "File f not found"

```

The local search succeeds only if a pointer to file f is in the local cache of s (line 1). Before returning the location of the file, the super-peer increases the priority of the corresponding cache item (line 2).

The global super-peer search protocol is as follows:

```

super-peer_search(s : super-peer,
    f : file_name)
1: perform a search in the super-peer
    network trying to locate a super-peer
    t which has an entry <f,q> in its cache
2: if (the search did not succeed)
3:   return ERROR "File f not found"
4: else
5:   insert <f,q> into s.F
6:   return <q,t>

```

The function `super-peer_search` is a wrapper for the search in the underlying super-peer network (line 1). Upon receipt of the search results, the pointer to the located file with the identity of the content provider peer q are added to the cache $s.F$ (line 5) and the pair $\langle q, t \rangle$ is returned (line 6).

2.4. Insert protocol

Our insert protocol is very simple. Each peer p once in a while sends information on the files which it possesses to one of the super-peers in its super-peer cache. This super-peer is selected randomly with a probability proportional to its priority in p 's super-peer cache.

2.5. Cache management

The super-peer and file caches are controlled differently. Whenever peer p receives positive feedback from super-peer s , the priority of s in cache $p.S$ is incremented by one which leads to the *in-cache least frequently used* (LFU) [7] cache management policy. The advantage of this method is the inherent memory property, which means that the priority of the super-peer is based on the amount of successful feedback it has provided in the past. The priority changes slowly, so one positive response from a completely unknown super-peer will not influence much the peer to super-peer assignment (this would be the case if we used one of the memoryless strategies [14] such as *least recently used* — LRU).

The management policy of the file caches should have two properties. First, it should enable the super-peers to adapt fast to the changing needs of their client peers. This is important, particularly in the initial stage of the super-peer lifetime, when it is contacted by random peers. Second, the cache management strategy should not ignore unpopular files. The cost of the search in the underlying super-peer network incurred by a cache miss is much higher for an unpopular file. In flooding protocols like Gnutella, looking for a file with only a few replicas in the whole network is much more expensive than searching for a popular file. The strategy proposed here is very simple. If the recently accessed file is in the cache, we increment its priority by one. Otherwise, we add this file with priority one higher than the highest priority of all cached items. This approach combines the fast adaptation of LRU with the memory property of LFU (in the remainder of the paper we call this caching method the *mixed* strategy). The high initial priority of the inserted item and the slow adaptation of the priorities of items in the cache prolong the caching period of less popular items.

3. Performance model

This section presents the method which we used for evaluating of the performance of the system described in Section 2. We start with a description of the semantic structure model. Then we analyze the properties of this model including the formal definition of the optimality of a particular arrangement of items in peer caches.

3.1. Model of the semantic structure

In our experiments we use a synthetic data model similar to the one introduced in [27]. This model assumes the existence of a number of (semantic) types for both files and peers labeled by $n \in \{1, \dots, N\}$, with N denoting the number of types. The semantic type of a file represents its content category (video, audio, software, etc.), while the type of a peer characterizes the content interest of the user. The number of files of type n is denoted by d_n , and the number of peers associated with type n is denoted by u_n .

Each peer periodically generates a request. The target file satisfying this request is selected randomly, according to a distribution that depends on the peer's type only. This distribution is specified by two parameters: the probability $p_n(m)$ that a request generated by a peer of type n will be targeted at a file of type m , and the probability $q_m(k)$ that a request for a file of type m will target the k -th file of this type. The distribution of $q_m(k)$ follows Zipf's distribution, which has been found to occur in real data traces [24]. The formulas for $p_n(m)$ can be written as follows:

$$\begin{cases} p_n(m) = \frac{1-\alpha}{m} / Z & , m = 1, \dots, N, m \neq n, \\ p_n(n) = (\alpha + \frac{1-\alpha}{n}) / Z & , \end{cases} \quad (1)$$

where Z is the normalizing constant chosen so that $\sum_m p_n(m)$ equals 1. More precisely, Z is given by $(1 - \alpha)H_N + \alpha$, where $H_N = \sum_{m=1}^N 1/m$ is the N -th harmonic number. Note that Z does not depend on n . The parameter α characterizes how strong the interest of users is for files of their own type.

Furthermore, we assume that our model satisfies two conditions. First, the numbers u_n of users of type n follow Zipf's law. Second, we require that the numbers of files d_n are all equal to M , independent of n . This second assumption may be controversial, but assuming that d_n follows Zipf's law as in [27] does not seem to be appropriate. This claim is based on our experience gained during our long-term measurements of BitTorrent, which is nowadays (March 2005) the most popular P2P network [20]. According to the statistics provided by one of the biggest `.torrent` distribution sites [6] in March 2005, the most popular content categories, which are movies and games, have much fewer items (5,000 movies and 2,000 games) than the relatively unpopular music (10,000 music files). In

this case the uniform distribution of d_n seems to be a better approximation of the real situation than Zipf's distribution, as it is at least independent from the popularity of individual categories.

It can be shown that the above assumptions imply that the popularity of file k of type m is proportional to $1/(mk)$, which is the same as in [27].

3.2. Optimal caching performance

Having a well-defined semantic structure of the data, we now analyze the caching performance for a given static assignment of peers to super-peers. We aim to identify such an optimal assignment, or in other words, the optimal arrangement of pointers to super-peers in the super-peer caches. It is generally not obvious how to define the *optimality* of a particular assignment. We describe the optimality in terms of performance and fairness stating that optimal means *Pareto optimal* [13]. An arrangement of items in super-peer caches is Pareto optimal if it is not possible to modify the contents of the cache of one peer in such a way that the fraction of requests produced by this peer that can be satisfied by the super-peers in its super-peer cache increases, while for all other peers this fraction does not decrease.

The above definition of the optimality does not say anything about the arrangements of items in the file caches. The reason for this is that the contents of the file caches are fully determined by the arrangements of items in the super-peer caches. To support this claim let's note that the content of a file cache of a super-peer depends only on the set of weak peers that send requests to this super-peer. However, a weak peer contacts a specific super-peer only if the pointer to this super-peer is in the weak peer's super-peer cache.

We can now prove that the semantic data model described in Section 3.1 has the following property.

Theorem A. *There exists an optimal arrangement of items in the super-peer caches such that the contents of the caches of all weak peers of the same semantic type are the same.*

PROOF. See the appendix.

This theorem can be shown to hold also in other models, including the one defined in [27].

Theorem A and the constructive proof thereof (see the appendix) can be used to show that the caching policies deployed locally at the peers guarantee a global convergence of the whole system to an optimal state. The policy of the super-peer caches tends to group similar peers under the same super-peers. As a consequence, the caches of the weak peers of one type are very similar, and in an optimal situation they are the same. Furthermore, weak peers constantly look for super-peers that give them better hit ratios, so every

non Pareto optimal assignment of peers-to-super-peers will be rearranged. The system will finally find the optimal state described by Theorem A.

4. Performance evaluation

This section presents the results of the experimental evaluation of our two-level caching architecture. We start with the explanation of the simulation parameters. Then we compare the performance of our system with the one offered by a system that uses only one level of semantic caches.

4.1. Experimental setup

We simulate a system consisting of 100,000 peers, 100 super-peers, 10,000 files, and 20 semantic types, so there are 500 files of the same semantic type. According to the SLYCK [5] portal (March 2005) a network simulated by us is comparable in size to some of the popular file sharing systems such as MP2P [4]. The value of the parameter α in Eq. 1 is set to 0.8. The size of the super-peer cache in the weak peers is set to 10, and the size of the file cache in the super-peers to 1,000. Before the simulation starts, all the super-peer caches have been filled with the identities of super-peers selected randomly and uniformly from the set of all super-peers. The file caches are initially empty. Each peer stores 50 files selected randomly, according to the distribution $p_n(m)$ described in Section 3.1.

The super-peers are organized into a Gnutella-like network. Simple request flooding in the super-peer network is employed for locating files which are not found in the local file caches. Note that this choice does not influence our results.

The simulation is performed in phases. At the beginning of each phase we select randomly and uniformly one of the peers. This peer generates a search request. Additionally, in every 1,000,000th phase, each of the 100,000 peers sends information about one of the files which it stores locally to a super-peer selected randomly from its super-peer cache (the file insert protocol). Note that the situation in which file inserts are performed by all peers simultaneously is the worst possible. During the whole experiment, 10,000,000 phases are performed. The first 1,000,000 phases are treated as the bootstrap. The statistics are collected starting from phase 1,000,001 on.

For comparison, we measure the performance of one-level caching. This reference system does not make use of super-peers. Similarly as in [27], each peer has a semantic cache of identities of peers that answered the peer's queries in the past. The caching policy used here is the same as the one deployed in the super-peer caches. The size of a peer's semantic cache is set to 300. This means that the total size of the caches used in the reference model,




Figure 2. The hit ratios of the two-level and one-level semantic caches.

which is 30,000,000 (100,000 peers with caches of size 300 each) is almost thirty times higher than the total size of the caches used in our super-peer architecture, which is 1,100,000 (100,000 super-peer caches of size 10 and 100 file caches of size 1,000).

4.2. Results of the experiments

In this section we assess different performance aspects of our super-peer network.

Figure 2 shows a comparison of the performance of our two-level caching architecture and the reference system which deploys one level of caches. For each file we compute the fraction of search requests targeting this file that is satisfied by one of the peer's direct neighbors. The direct neighbors in the reference model are the nodes stored in the peer's cache, while in the two-level scheme they are defined as the super-peers contained in the super-peer cache. The files are sorted according to their semantic type. Each file is assigned a unique identifier which we call its file rank — file k of type m has rank $(m-1)M+k$. The top left subplot of Figure 2 shows the hit ratios observed for our two-level caching architecture. The top right plot depicts results of the same experiment performed in a system with only one level of semantic caches. The bottom subplot presents a comparison of the efficiency of both caching schemes; it was obtained by applying curve fitting based on nonlinear regression [2] to the points shown in the top two plots.

In most of the cases the super-peer algorithm outperforms the one-level design, even though the total number of items cached in the two-level system is much lower than in the reference model. Only for a small number of the most popular files the two approaches achieve comparable results. The average cache hit ratio defined as the percentage of all requests that were satisfied from the cache during the whole simulation for our super-peer architecture was 71% against 56% for the one-level caching scheme.

In order to investigate the properties of the mixed caching method we repeat the series of experiments described in Section 4.1 for the LRU and LFU policies applied to the file caches. The results of the comparison are presented in Figure 3 for the files in the ten most (least) popular categories at the top (bottom). Also in this case we performed curve fitting to improve the clarity of the plot.

The LFU strategy promotes the most popular files, but neglects the unpopular content, achieving cache hit ratios comparable with these observed for the LRU. The mixed method, while still providing high cache hit rates for the popular items, performs much better for the unpopular content. We computed also the average cache hit ratios for LFU, LRU, and the mixed strategy, which are 62%, 63%, and 71%, respectively.

According to the protocol presented in Section 2.3, if the requested file is not cached by one of the super-peers known to the request initiator, a search in the underlying super-peer network is performed. It is possible that some of the files




Figure 3. A comparison of different caching strategies applied to file caches.

possessed by the weak peers are not cached by any of the super-peers, and as a consequence, cannot be found during the super-peer search. We investigate the influence of different file cache management strategies on the number of requests that can not be satisfied because the located file is not indexed by any super-peer. Figure 4 shows the fraction of requests addressing files of a particular type that could not have been satisfied because the located file was not indexed by any super-peer (the smaller the better).

It should be noticed that the fraction of missed requests is generally very low. We can clearly see that the mixed strategy in most of the cases outperforms the other two. The strong preference of the most popular files exhibited by the LFU policy is visible in the statistics for types 5 and 6. LRU is more fair than LFU for less popular files, but still falls behind the mixed strategy.

5. Conclusions and future work

In this paper we have introduced semantic caches deployed on top of the super-peers architecture to improve the performance and scalability of the search in a P2P network. Our system does not make any assumptions according to the type of the content shared in the network, and does not require users to explicitly specify their interests or manually categorize the content.

The semantic caching model for a content sharing P2P

network proposed in this paper is based on, and derives its advantages from three key issues. First, we use the heterogeneity of the nodes by placing the semantic caches of file pointers only at the super-peers, which are selected nodes with higher capacities. These file caches are used by multiple peers at the same time, which shortens the cache warm-up period and increases the utilization of the cache. Second, weak peers maintain caches of super-peers which are most likely to know the location of the files they are interested in. Third, we introduce the mixed cache-management policy for the file caches that achieves much higher cache hit ratios for less popular files while still being very efficient for the most popular items. Furthermore, the probability that a file is not known to any super-peer in the network is much lower for the mixed policy than for LRU or LFU.

The analysis and the experimental evaluation on the synthetic data set give us an idea about the performance gain of the proposed infrastructure. However, in order to determine how a super-peer architecture based on semantic caches performs in a real environment, more elaborate experiments are needed. For this purpose, we plan to use our 2-year traces of the `suprnova.org` website which contains popularity statistics of the content distributed in the BitTorrent P2P network [20]. We will also consider extensions of the cache management algorithms. Interesting results may be achieved by allowing peers to adapt not only content but also the sizes of the caches.




Figure 4. The file coverage of the super-peer caches.

References

- [1] <http://gnutella.wego.com>.
- [2] <http://www.curvefitting.com>.
- [3] <http://www.kazaa.com>.
- [4] <http://www.mp2p.net>.
- [5] <http://www.slyck.com>.
- [6] <http://www.thepiratebay.org>.
- [7] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web caching and zipf-like distributions: Evidence and implications. In *INFOCOM*, pages 126–134, March 1999.
- [8] Y. Chawathe, S. Ratnasamy, L. Breslau, and S. Shenker. Making gnutella-like p2p systems scalable. In *Proceedings of the ACM SIGCOMM 2003 Conference, Karlsruhe, Germany*, August 2003.
- [9] A. Crespo and H. Garcia-Molina. Routing indices for peer-to-peer systems. In *Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCS'02)*, Vienna, Austria, July 2002.
- [10] F. L. Fessant, S. Handurukande, A.-M. Kermarrec, and L. Massoulié. Clustering in peer-to-peer file sharing workloads. In *3rd International Workshop on Peer-to-Peer Systems (IPTPS)*, San Diego, CA, February 2004.
- [11] A. Ganesh, A.-M. Kermarrec, and L. Massoulié. Peer-to-peer membership management for gossip-based protocols. *IEEE Transactions on Computers*, 52(2), February 2003.
- [12] S. Handurukande, A.-M. Kermarrec, F. L. Fessant, and L. Massoulié. Exploiting semantic clustering in the edonkey p2p network. In *11th ACM SIGOPS European Workshop (SIGOPS)*, Leuven, Belgium, September 2004.
- [13] M. Hutter. Self-optimizing and Pareto-optimal policies in general environments based on Bayes-mixtures. In *Proceedings of the 15th Annual Conference on Computational Learning Theory (COLT 2002)*, Lecture Notes in Artificial Intelligence. Springer, 2002.
- [14] N. Laoutaris, S. Syntila, and I. Stavrakakis. Meta algorithms for hierarchical web caches. In *IEEE International Performance Computing and Communications Conference (IEEE IPCCC)*, Phoenix, Arizona, April 2004.
- [15] X. Li and J. Wu. *Handbook of Theoretical and Algorithmic Aspects of Ad Hoc, Sensor, and Peer-to-Peer Networks*, chapter Searching Techniques in Peer-to-Peer Networks. CRC Press, 2005.
- [16] A. Loser, M. Wolpers, W. Siberski, and W. Nejdl. Semantic overlay clusters within super-peer networks. In *International Workshop on Databases, Information Systems, and P2P Computing, collocated with 29th International Conference on Very Large Databases (VLDB2003)*, Berlin, Germany, september 2003.
- [17] D. S. Milojicic, V. Kalogeraki, R. Lukose, K. Nagaraja, J. Pruyne, B. Richard, S. Rollins, and Z. Xu. Peer-to-peer computing. Technical Report HPL-2002-57, HP, March 2002.
- [18] W. Nejdl, M. Wolpers, W. Siberski, C. Schmitz, M. Schlosser, I. Brunkhorst, and A. Loser. Super-peer-based routing and clustering strategies for rdf-based peer-to-peer networks. In *12th International World Wide Web Conference (WWW2003)*, Budapest, Hungary, may 2003.
- [19] A. Parker. The true picture of peer-to-peer filesharing, 2004. <http://www.cachelogic.com/>.
- [20] J. Pouwelse, P. Garbacki, D. Epema, and H. Sips. The bittorrent p2p file-sharing system: Measurements and analysis. In *Proceedings of the 4th International Workshop on Peer-To-Peer Systems (IPTPS'05)*, Ithaca, New York, february 2005.
- [21] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content addressable network. In *Proceedings of ACM SIGCOMM 2001*, 2001.
- [22] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. *Lecture Notes in Computer Science*, 2218:329–350, 2001.
- [23] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy. An analysis of internet content delivery systems. In *Proceedings of the 5th Symposium on Operating Systems Design and Implementation (OSDI)*, Boston, MA, December 2002.
- [24] K. Sripanidkulchai, B. Maggs, and H. Zhang. Efficient content location using interest-based locality in peer-to-peer systems. In *INFOCOM*, 2003.

- [25] I. Stoica, R. Morris, D. Karger, and M. F. K. H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the 2001 conference on applications, technologies, architectures, and protocols for computer communications*, pages 149–160. ACM Press, 2001.
- [26] S. Vanichpun and A. M. Makowski. The output of a cache under the independent reference model — where did the locality of reference go? In *SIGMETRICS 2004/PERFORMANCE 2004: Proceedings of the joint international conference on Measurement and modeling of computer systems*, pages 295–306, New York, 2004. ACM Press.
- [27] S. Voulgaris, A.-M. Kermarrec, L. Massoulie, and M. van Steen. Exploiting semantic proximity in peer-to-peer content searching. In *10th IEEE Int'l Workshop on Future Trends in Distributed Computing Systems (FTDCS 2004)*, Suzhou, China, may 2004.
- [28] B. Yang and H. Garcia-Molina. Designing a super-peer network. In *IEEE International Conference on Data Engineering*, Bangalore, India, March 2003.
- [29] B. Y. Zhao, L. Huang, J. Stribling, S. C. Rhea, A. D. Joseph, and J. D. Kubiatowicz. Tapestry: A resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications*, 22(1):41–53, January 2004.

Appendix

PROOF OF THEOREM A. Consider the file cache $s.F$ of super-peer s , and denote by $r_n(s)$ the fraction of all requests submitted to s that are issued by peers of type n . Assuming the Independent Reference Model [26] of requests and a reasonable caching policy at the file caches, the cache hit ratio of a peer of type n is non-decreasing function of $r_n(s)$ — the more requests are produced by peers of type n , the better cache $s.F$ adapts to the needs of peers of this type.

We say that an arrangement of items in the super-peer caches is *structured* if the caches of all peers of the same semantic type are the same. Consider a non-optimal structured arrangement of items in the super-peer caches. Non-optimal means (see Section 3.2) that it is possible to modify a super-peer cache of one of the peers, say p , in such a way that two conditions hold. First, the fraction of requests produced by p that can be satisfied by the super-peers in p 's super-peer cache increases. Second, for all other peers this fraction does not decrease.

A cache modification for a peer p of type n influences the values of the $r_n(s)$ of some of the file caches at the super-peers. Removing any super-peer s from p 's cache decreases (or does not influence) $r_n(s)$, and increases (or does not influence) $r_m(s)$, where m is different from n . Subsequently adding any super-peer s to p 's cache has the opposite effect — $r_n(s)$ increases (or does not change), and $r_m(s)$ decreases (or does not change). Furthermore, if for some l , $r_l(s)$ increased (decreased) after modifying p 's cache, then it will also increase (decrease) when we perform the same

modification to the cache in p' , where p' is a peer of the same type as p . We showed earlier that the cache hit ratio of a peer of type l at super-peer s is a monotonic function of $r_l(s)$. We conclude that if a certain modification in a peer's super-peer cache improves the hit ratios at some peers and does not decrease the hit ratios at all other peers, then applying the same modification to the cache in a peer of the same type will improve the same hit ratios, and will not decrease the other.

We have just shown that if a certain modification of a super-peer cache of peer p improves the non-optimal, structured arrangement of super-peer caches, then by applying the same modification to the super-peer caches of all peers of the same type as p the arrangement can be improved by at least the same amount. Such a modification results in an arrangement that is again structured. There is, however, a finite number of (structured) arrangements, which means that we cannot endlessly improve. At some point we will end up with a structured arrangement which is optimal.