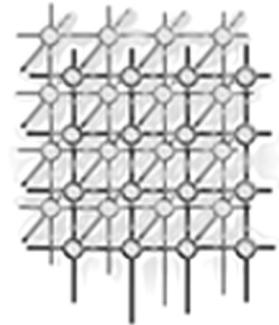


TRIBLER: a social-based peer-to-peer system

J. A. Pouwelse¹, P. Garbacki¹, J. Wang¹, A. Bakker², J. Yang¹,
A. Iosup¹, D. H. J. Epema^{1,*},[†], M. Reinders¹, M. R. van Steen²
and H. J. Sips¹



¹Department of Computer Science, Delft University of Technology, The Netherlands
²Department of Computer Science, Vrije Universiteit, The Netherlands

SUMMARY

Most current peer-to-peer (P2P) file-sharing systems treat their users as anonymous, unrelated entities, and completely disregard any social relationships between them. However, social phenomena such as friendship and the existence of communities of users with similar tastes or interests may well be exploited in such systems in order to increase their usability and performance. In this paper we present a novel social-based P2P file-sharing paradigm that exploits social phenomena by maintaining social networks and using these in content discovery, content recommendation, and downloading. Based on this paradigm's main concepts such as taste buddies and friends, we have designed and implemented the TRIBLER P2P file-sharing system as a set of extensions to BitTorrent. We present and discuss the design of TRIBLER, and we show evidence that TRIBLER enables fast content discovery and recommendation at a low additional overhead, and a significant improvement in download performance. Copyright © 2007 John Wiley & Sons, Ltd.

Received 24 September 2006; Revised 25 January 2007; Accepted 30 January 2007

KEY WORDS: peer-to-peer; social-based; taste buddies

1. INTRODUCTION

Traditional peer-to-peer (P2P) file-sharing systems focus exclusively on technical issues and are therefore unable to leverage the power of social phenomena. However, we believe that social phenomena such as friendship, trust, and a sense of community may be at least as important as technical issues, and may indeed have a large positive impact on the usability and performance of P2P file-sharing and content-delivery systems. For example, viewing users as *social partners* rather than as

*Correspondence to: D. H. J. Epema, Department of Computer Science, Delft University of Technology, P.O. Box 5031, 2600 CD, Delft, The Netherlands.

[†]E-mail: d.h.j.epema@tudelft.nl



solitary *rational agents* [1] could alleviate the problem of free-riding [2] by exploiting the fact that people tend not to steal (bandwidth) from the social group they belong to.

To confirm our belief, we propose in this work a novel *social-based P2P file-sharing paradigm*, which facilitates the formation and maintenance of social networks and exploits social phenomena for improved content discovery, recommendation, and sharing. Our contribution is threefold. First, we relate the social-based P2P file-sharing paradigm to the current research challenges in P2P research (Section 2). Second, we present the design and implementation of TRIBLER, which adheres to the social-based paradigm by adding social-based functionality to the widely used BitTorrent system (Section 3). To facilitate the formation and maintenance of social networks, TRIBLER introduces permanent user identifiers, and will in the future be able to import existing user contacts from other social networks such as MSN (Section 4). Rather than using direct content-based searching, TRIBLER performs content discovery and recommendation based on the notion of *taste buddies*, that is, users with the same tastes or interests (Section 5). Third, we show evidence that TRIBLER achieves a significant improvement in download performance, by invoking the joint efforts of social peer groups (Section 6). In Section 7 we discuss related work, and in Section 8 we present our conclusions and ideas for future work. The full TRIBLER documentation and source code are available from <http://Tribler.org>.

2. RESEARCH CHALLENGES IN P2P FILE SHARING

With current P2P file-sharing systems continuously having more than 1 000 000 users, their performance and behavior have become of great interest. Starting in 2003, we have studied the performance of BitTorrent [3], which has for a number of years been the most popular P2P file-sharing system. Based on this work and on related studies, we formulate the following five grand research challenges for P2P file sharing, and we argue for the importance of the social-based paradigm in solving these challenges. In particular, with our social-based P2P network called TRIBLER we want to address all these challenges.

The most difficult research challenge is the *decentralization* of the functionality of a P2P system across its peers. Full decentralization eliminates the need for central elements in the system, which must be set up and maintained by some party and which may form serious bottlenecks, points of failure, or security threats. In particular, connecting to the network and validating user identities are difficult to implement without any central element. To date, no P2P file-sharing system exists which fully decentralizes all functionality efficiently and without a high risk of a loss of integrity. BitTorrent is not fully decentralized as it depends on Web sites and trackers for finding content. Social groups form a natural method to efficiently decentralize P2P systems because people/peers who know each other tend to exchange information.

The second challenge is to guarantee the *availability* of a P2P system as a whole. The operation of such a system should not depend on the availability of any particular participating peer, or of any central component (something we do not want if we aim for decentralization), as the failure of such a component can be disruptive to service [3]. Given the short periods of availability of peers (in [3] we found less than 4% of the peers to have an uptime of over 10 hours), the availability problem is critical. Proven social incentives such as rewards and social recognition could stimulate users to leave their P2P software running for longer periods, thus improving the overall availability of the network.

The third challenge is to maintain the *integrity* of the system and to achieve *trust* among peers. By definition, P2P systems use donated resources. However, donors cannot always be trusted, and



maintaining system integrity has proven to be difficult in operational P2P systems [4]. Data at several levels can be attacked in a P2P system, namely system information (e.g., pointers to content), metadata, and the actual content itself. This significant problem, often ignored by P2P system designers, can be solved with a social-based network, in which users actively help to clean or remove polluted data and users can select trustworthy representatives.

The performance of a P2P system highly depends on peers donating resources. Even though the resource economy is by definition balanced (e.g., every Mbyte downloaded corresponds to a Mbyte uploaded), autonomous peers are free to decide whether to donate resources or not. Hence, as the fourth challenge, *providing proper incentives* is vital to induce cooperation and to achieve good performance [2]. Again, social recognition can help to alleviate this problem.

The fifth challenge in P2P systems is to achieve *network transparency* by solving the problems caused by dynamic IP addresses, NAT boxes, and firewalls. The Internet has fundamentally changed owing to the wide-spread use of these three technologies, and, as a consequence, peers no longer have the freedom to send anything anywhere without the help of another peer acting as a mediator. Social networks enable communicating peers to automatically select trusted mediators from the members of their social proximity that are online, thus eliminating the need for fixed mediators.

3. THE ARCHITECTURE OF TRIBLER

In this section, we present the architecture of our TRIBLER social-based P2P file-sharing system, which is built on top of the BitTorrent protocol. Figure 1 depicts this architecture, with rectangles representing its modules and files, and extrusions representing *makes-use-of* relationships. To achieve backwards compatibility with the existing BitTorrent P2P system, we have only made modifications and extensions to the existing BitTorrent client software. We have based our system on the popular ABC open-source BitTorrent client [5] so as to have a tested code base for our implementation and a large user base in a relatively short period of time. We now discuss several important concepts in TRIBLER.

3.1. Social groups

The prime social phenomenon that we exploit in TRIBLER is that ‘kinship fosters cooperation’ [6]. In other words, a similar taste for content can form the foundation for an online community with altruistic behavior. In order to implement effective social groups in TRIBLER, we use an approach borrowed from evolutionary biology (see, for instance, [6]): we have implemented the ability to distinguish friend, foe (e.g., a polluter), and newcomer. For this, we de-anonymize peers and facilitate social-group formation. De-anonymization is achieved by having every user choose a nickname; TRIBLER transfers user nicknames between users automatically. The *Social Networking* module in Figure 1 is responsible for storing and providing information regarding social groups (the group members, their recently used IP numbers, etc.).

3.2. Megacaches

Virtually all current P2P file-sharing systems lack persistent ‘memory’ about previous activity in the network; peers usually exchange queries for files and file metadata, but completely ignore other types

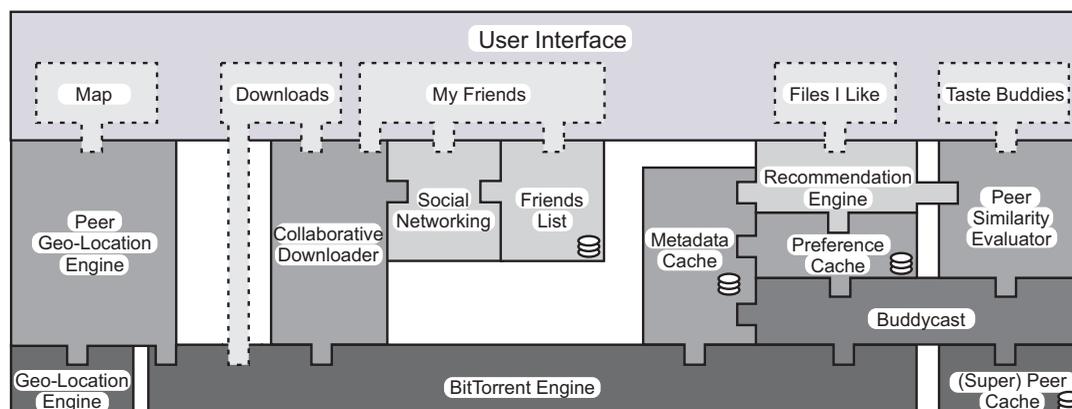


Figure 1. The system architecture of TRIBLER.

of information. The *context information* that needs to be saved in order to improve the performance consists of information on social relations, altruism levels, peer uptimes, taste similarities, etc. In TRIBLER, every piece of context information received by a peer that is relevant to it based on its interests and tastes is stored locally in its so-called *Megacaches*, and this information is exchanged within social groups using an epidemic protocol [7] (the Buddycast protocol, see Section 5). The small database icons in Figure 1 identify the four Megacaches in TRIBLER, which are the *Friends List* with information on social networks, the *(Super) Peer Cache* with information on super-peers and peers in general, the *Metadata Cache* with file metadata, and the *Preference Cache* with the preference lists of other peers. The sizes of the *Friends List*, the *(Super) Peer Cache*, and the *Preference Cache* are below 10 MB at any time, which for the *Preference Cache* is enough to store the preference lists of thousands of taste buddies (see also Section 5).

The main problem concerning the Megacaches is the overhead traffic required to keep them up-to-date. For the *Metadata Cache*, we have observed that in BitTorrent the number of newly injected files per day is limited to roughly 1500 [3] when *content pollution* [4] is kept to a minimum. Then, by reducing the average size of the metadata for each file to just 400 bytes using Merkle hashes [8], the overhead is reduced to approximately 600 KB per day. With this amount of overhead, all metadata can be replicated among all peers, moving content discovery from network-based keyword searching to local metadata browsing.

3.3. Taste-buddy-based content discovery

Locating content is critical for P2P systems. Current solutions are based on one or a combination of query flooding, distributed hash tables, and semantic clustering. We take a next step by connecting *people* with similar tastes called *taste buddies* instead of focusing on *files*, and by using full metadata replication.



Using the *Files I like* module (see Figure 1), each peer indicates its preference for certain files expressed as a number between 1 and 5. By default, the preference list of a peer is filled with its most recent downloads. We have developed an algorithm called *Buddycast* which uses an epidemic protocol to exchange preference lists using the overlay swarm (see Section 3.6) and which can efficiently discover a user's taste buddies (see Section 5). The *Peer Similarity Evaluator* module in Figure 1 is able to compare preference lists and determine the similarity in taste of two peers, which is measured as the cosine of their vectors of user ratings of content.

The *Recommendation Engine* module is able to compile a list of files a user most likely wants. First, a user-item rating matrix is built from the preference lists [9], and then a user-based recommendation is generated by TRIBLER, based on standard collaborative filtering techniques. Finally, the user interface facilitates metadata browsing by augmenting each file entry with the estimated interest to the user.

3.4. Downloading

The *BitTorrent Engine* module in Figure 1 downloads files using a BitTorrent-compatible protocol. This module can also use the *Collaborative Downloader* module's capabilities to achieve a significant increase in file download speed by exploiting idle upload capacity of online friends (see Section 6).

3.5. User interface

The user interface is key to making a social-based network usable and, as such, is a critical part of TRIBLER. The *User Interface* module from Figure 1 is split into five components: *Map*, *Downloads*, *My Friends*, *Files I like*, and *Taste Buddies*. The use of the *Downloads*, *Files I like*, and *Taste Buddies* modules is clear from the description above. The main goal of the interface is to facilitate the formation of social groups. For this purpose, the *My Friends* module clearly displays the friends, the friends-of-friends, and the taste buddies of the user. This *visual proximity* gives the user a more personal contact with his peers, and may help reduce asocial behavior.

Another goal of the user interface design was to ease the process of visual identification of potential collaborators. When a user is downloading a file, the observed IP addresses of members of the corresponding download swarm are geo-located and then displayed on a world map using the *Map* module. We have built a *Peer Geo-Location Engine* module on top of a freely available *Geo-Location Engine* (<http://hostip.info>). The user interface also facilitates the use of the *Collaborative Downloader* module: the user can see which friends helped them in the past, which friends they donated bandwidth to, and which friends who are currently online can help speedup their new downloads.

3.6. Bootstrapping

Finding other peers in a P2P system after software installation is called bootstrapping. In BitTorrent, peers have to repeatedly connect to a tracker in order to discover other peers. Furthermore, the original BitTorrent protocol restricts communication to within the swarms of peers that download the same files, making the bootstrapping process unnecessarily repetitive. To solve the bootstrapping problem in TRIBLER, we use two mechanisms. First, a TRIBLER peer automatically contacts one of a set of pre-known *super-peers* only once immediately after installation in order to obtain an initial list of other peers in the system (through a Buddycast message, see Section 5), so that it can start participating



Table I. Performance of the TRIBLER system.

Test description	Performance (ops s ⁻¹)
Geo-lookup	1730
Overlay swarm connect	7045
Connect+challenge/response	865
Connect+challenge+Buddycast	844

in the epidemic information dissemination in TRIBLER. Second, we define a special *overlay swarm*, which is a swarm of which every TRIBLER peer is a member that has no tracker and that is used for content and peer discovery, again using Buddycast.

3.7. Performance of TRIBLER

In order to assess the overhead incurred by some of the operations in the TRIBLER protocol, we have carried out some performance tests on a powerful computer (a 4-CPU 2.0 GHz machine with 16 GB of main memory and a 1 Gbps Internet connection). Table I shows the results of these tests. The first test determines the number of geo-lookups this machine can handle per second. The second test shows the performance of connecting to a peer in the overlay swarm with the standard BitTorrent protocol messages. In the third test, in addition also the public key of the peer that connects is validated using a challenge/response algorithm (see Section 4). The fourth test adds to the third test the exchange of a preference list with 100 file names by means of the Buddycast algorithm. These tests show that a single peer can handle a significant workload, although indeed a certain amount of overhead is added to the BitTorrent protocol.

4. SOCIAL NETWORKING

A fundamental limitation in most file-sharing systems is the *session boundary*—all context information is lost when a user disconnects from the network. As a result of dynamic IP numbers, it is difficult to store context information about peers across sessions. Storing long-term context information in databases such as our Megacaches enables the existence of trust-based social groups, but only if the user identities are stable. To solve this problem, we have introduced permanent, unique, and secure peer identifiers (*PermIDs*), which are the public keys of a public-key scheme using elliptic-curve cryptography, and which are exchanged by e-mail. To prevent peers from faking the identity belonging to a PermID (*spoofing*), we have implemented a challenge-response mechanism for validating PermIDs. In this mechanism, a random number selected by the challenger is encrypted by the challengee with its private key and subsequently decrypted by the challenger with the corresponding public key—authentication succeeds when the result is identical to the original number. Social-network creation in TRIBLER will be facilitated by the ability to import contacts from other social networks of which peers are members, such as MSN and GMail.

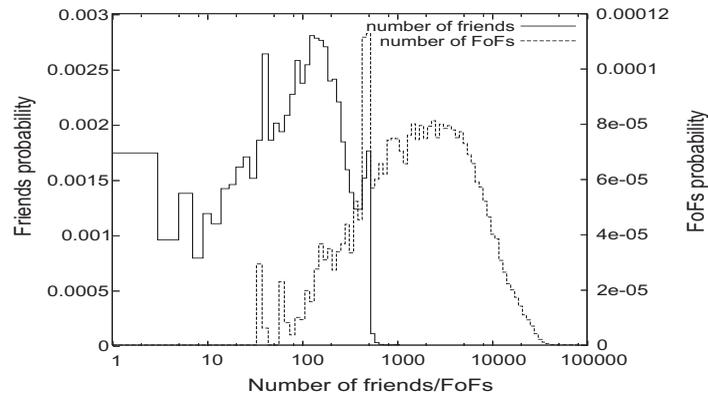


Figure 2. The probability density of the numbers of friends and friends-of-friends (FoFs) in Friendster.com.

We will use Bloom filters [10], which are very dense hash-table-like data structures for storing and (probabilistically) testing set membership, for distributing and pairwise comparing the contents of the Megacaches. Because of their reduced size, Bloom filters can significantly reduce the bandwidth requirements of epidemic protocols, which are the basis of our solution for content discovery and social networking. In a recent prototype [11] of a design for decentralized swarm discovery in TRIBLER, we have already implemented Bloom filters for filtering peers from messages that are already known to their destinations.

The size of a Bloom filter depends on the number of expected connections peers have. To assess their size while we do not yet have sufficient information on the operation of TRIBLER, we resort here to analyzing another existing social network, Friendster.com. We have created a crawler for this network and we have extracted 3.3 million relations between 27 000 people. Figure 2 shows the probability density of the numbers of friends and friends-of-friends in Friendster.com. In this network, we find that a person has on average 243 friends and 9147 friends-of-friends. These figures are to within an order of magnitude similar to the figures reported in [12] for several P2P file-sharing networks. Based on these numbers, we can compute that 260 bytes are needed to discover the common friends-of-friends of two peers using a Bloom filter. This very low bandwidth requirement will enable TRIBLER peers to exchange information with *thousands* of other peers in a short time frame, which is a significant improvement over traditional epidemic protocols.

5. THE BUDDYCAST ALGORITHM

In order to make recommendations and to enable peer and content discovery, and, more generally, to disseminate the contents of the Megacaches in the TRIBLER overlay swarm, we have designed the *Buddycast algorithm*, which is an epidemic protocol that works as follows. Each peer maintains in its Megacache a number of taste buddies with their preference lists and a number of random

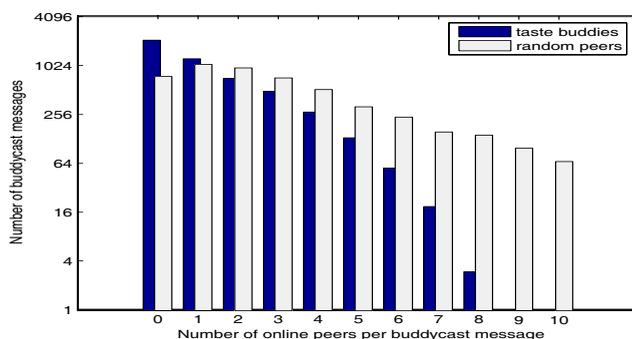


Figure 3. The availability of peers in Buddycast messages.

peers (i.e. peers for which no preference list is known). Periodically (in the current version of TRIBLER every 15 seconds), a peer either connects to (a) one of its taste buddies (*exploitation*) or to (b) a random peer (*exploration*) to exchange a *Buddycast* message, or it replies to such a message received from another peer. In contrast to other epidemic protocols such as Newscast [7], we use both exploitation and exploration, we limit the randomness of peer selection during the exploitation, and we implicitly cluster similar peers into (trusted) social groups. This is very close to the approach proposed in [13].

A Buddycast message contains the identities of a number (currently 10) of taste buddies along with their top-10 preference lists, a number (currently also 10) of random (and fresh, see below) peers, and the top-50 preferences of the sending peer. The age of each peer is included in the message to help others know the freshness of peers. After exchanging Buddycast messages, both peers merge the information in the message received into their own Megacache. To maximize the exploration and avoid redundant messages, every peer also maintains a list with the most recently contacted random peers, and avoids reconnecting to the peers in this list for a certain period of time (currently 4 hours). Users can disable the functionality of Buddycast to protect their privacy.

To find a good balance between exploration and exploitation, an *exploration-to-exploitation ratio* λ for selecting either a random peer or a taste buddy is used, which is currently set to 1, but which we may want to make adaptive. In the case of exploration, a random peer is selected based on its freshness, to improve the chance of its still being online when connecting to it. In the case of exploitation, a taste buddy is selected from a number (currently 100) of the most recently contacted taste buddies based on its similarity with the sending peer. As a consequence, the fresher or more similar a peer, the higher its chance to be selected.

For a first assessment of the operation of the Buddycast algorithm, we have investigated the availability of the peers contained in Buddycast messages at the moment of receiving such messages. To this end, we have run a continuously fresh TRIBLER peer for a period of 519.7 hours, and we have probed all the peers in every Buddycast message it received to see how many of them were still online. At the time of this experiment, there were about 15 000 TRIBLER clients in the world who had actually downloaded files. The results are shown in Figure 3 (note the logarithmic scale on the vertical axis).

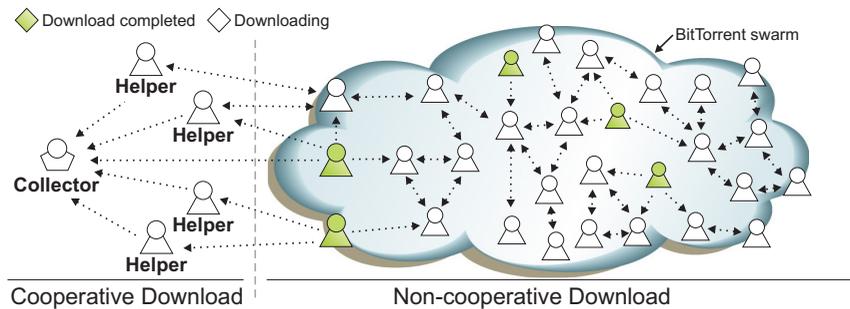


Figure 4. Overview of collaborative downloading.

During the experiment, our peer received 5049 messages in total. As can be seen in the figure, often at most only four taste buddies in the Buddycast messages were still online, which agrees with our earlier finding that online periods in BitTorrent are short [3]. From the figure we also see that random peers have a higher chance to be online, which indicates that selecting peers based on freshness can benefit the effectiveness of exploration.

6. COLLABORATIVE DOWNLOADING

In this section we present the *collaborative download protocol* called 2Fast that we have developed for TRIBLER, in which a user invokes the help of his friends to speed up downloads. Early downloading protocols (e.g., Gnutella) have no incentives for donating upload bandwidth, which has serious limitations in real environments, because unconstrained bandwidth sharing is sensitive to free-riding [2]. The BitTorrent Tit-for-Tat mechanism was the first system which offered an incentive for uploading. The current BitTorrent mechanism also has its disadvantages, because without enough seeding peers (i.e., peers which possess the complete file in question), the download speed of a peer is restricted by BitTorrent's Tit-for-Tat bartering protocol to its upload link capacity. Hence, peers with asymmetric Internet access, such as ADSL or ADSL-2, cannot fully use their download capacity.

The 2Fast protocol makes use of social groups, where members who trust each other collaborate to improve their download performance (see Figure 4). The idea of downloading with the help of others was first introduced in [14], where altruistic peers contribute their bandwidth by joining a swarm even if they are not interested in the content being distributed in this swarm. The inherent assumption of sufficient altruism in the network without any incentives makes this simple approach impractical in real-world environments. Our 2Fast protocol solves this problem by introducing social-group incentives. For a full account of 2Fast, see [15].

In 2Fast, peers from a social group that decide to participate in a collaborative download take one of two roles: they are either *collectors* or *helpers*. A collector is the peer that is interested in obtaining a complete copy of a particular file, and a helper is a peer that is recruited by a collector to assist in downloading that file. Both the collector and the helpers start downloading the file using the classical

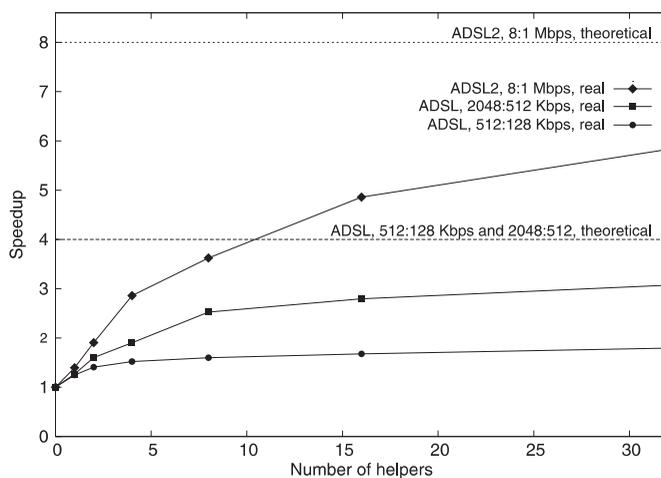


Figure 5. The impact of the number of helpers on the download speedup.

BitTorrent Tit-for-Tat and the collaborative download extensions. As in BitTorrent, a helper selects a chunk of the file for downloading based on the rarest-first policy among the chunks in possession of its bartering partners. However, before actually downloading this chunk, it asks the collector for approval, which will only be granted when the chunk is unique, that is, when no other helper already downloads the same chunk. After downloading a file chunk, the helper sends the chunk to the collector without requesting anything in return. In addition to receiving file chunks from its helpers, the collector also optimizes its download performance by dynamically selecting the best available data source from the set of helpers and other peers in the BitTorrent network using the default mechanisms of BitTorrent, which prefer peers that upload at higher rates. As helpers give priority to collector requests, they are preferred as data sources.

We have implemented and tested 2Fast in a real environment. For this we have selected a middle-sized swarm of around 1900 peers with only 6% seeds, distributing a 1.2 GB file. These numbers remained almost unchanged during our experiments. We have performed tests for three download-to-upload bandwidth ratios from standard Internet package offerings: low-end ADSL with a ratio of 512:128 Kbps, high-end ADSL with a ratio of 2048:512 Kbps, and ADSL-2 with a ratio of 8:1 Mbps. It should be noted that we could only impose these bandwidths on the collector and the helpers, which were under our control, but not on the peers in the Internet that they selected as their bartering partners.

As a performance metric of our system, we use the ratio of the download time achieved by a peer obtaining a file all by itself to the corresponding time for a collaborative group (*speedup*). The theoretical maximum speedup is achieved when a peer can fully use its download bandwidth, and so it is equal to the ratio between the download and upload link capacities. Thus, for ADSL and ADSL-2 the maximum achievable speedup is 4 and 8, respectively.

Figure 5 shows the obtained speedups for the number of helpers in the range from 0 to 32. The total download time was decreased by a factor of almost two for low-end ADSL, of more than three for



high-end ADSL, and of almost six for ADSL-2. The difference between the theoretical and achieved speedups is mainly due to the influence of the seeders (when there are many seeders to begin with, the potential speedup is restricted) and of the delays experienced by helpers when requesting unique file chunks from peers. The more helpers that are involved, the more restrictive the unique file chunk selection method is, and, consequently, the longer the time needed to find a bartering partner for such a chunk. This time is further increased in the case of low-end ADSL by the fact that then the collector and the helpers were discriminated as peers with an upload bandwidth which is below average [3].

7. RELATED WORK

The idea of exploiting the natural social connections between the humans behind the computers in large-scale P2P networks is starting to become a major research topic. To date, methods based on social clustering were applied in P2P networks to limited aspects of content distribution [16], user communities formation [17], and collaborative service provisioning [9]. In [16], a system which uses knowledge discovery techniques for overlay network creation is presented. By automatically clustering users based on their preferences, the system enables content location and improves the performance of content sharing. In [17], a simple general-purpose system is proposed which groups peers based on the similarity of their keyword searches. The authors give evidence on how their system can be used to form and maintain communities of users. An extensive experimental analysis of several collaborative filtering methods is given in [9]. TRIBLER is the first system which exploits social phenomena to address all the research challenges in P2P file-sharing networks mentioned in Section 2. The BTSlave project [18] extends the BitTorrent protocol by introducing special-purpose peers called *repeaters* that use their idle bandwidth to increase content replication in the swarm. In contrast to TRIBLER, in BTSlave a peer is not aware of the help it is receiving, and so it cannot coordinate its helpers as the collector in 2Fast does. Owing to the lack of coordination, BTSlave repeaters can end up competing for the same content instead of collaborating.

8. CONCLUSION AND FUTURE WORK

In this paper we have presented a novel paradigm for the design of P2P file-sharing networks based on social phenomena such as friendship and trust. Following the paradigm's main concepts of taste buddies and friends, we have designed and implemented the TRIBLER P2P file-sharing system. We have described how TRIBLER can help to automatically build a robust semantic and social overlay on top of BitTorrent, one of the most popular P2P file-sharing systems. We have shown how various TRIBLER components can yield good performance with respect to existing solutions. In particular, we have presented evidence that collaborative downloading yields a significant speedup when used in a real BitTorrent swarm. Last but not least, we have shown how with TRIBLER we have made a start in addressing the five major P2P research challenges introduced in Section 2.

As future work, we intend to extend TRIBLER with a reputation system, with tag-based content navigation, with video-on-demand, and with application-level multicasting for video streaming. In addition, we are planning a much more detailed analysis of the Buddycast algorithm.



REFERENCES

1. Nielson SJ, Crosby SA, Wallach DS. A taxonomy of rational attacks. *Proceedings of the 4th International Workshop on Peer-to-Peer Systems (IPTPS) (Lecture Notes in Computer Science, vol. 3640)*. Springer: Berlin, 2005; 36–46.
2. Adar E, Huberman BA. Free riding on Gnutella. *Technical Report*, Xerox PARC, August 2000.
3. Pouwelse JA, Garbacki P, Epema DHJ, Sips HJ. The BitTorrent P2P file-sharing system: Measurements and analysis. *Proceedings of the 4th International Workshop on Peer-to-Peer Systems (IPTPS) (Lecture Notes in Computer Science, vol. 3640)*. Springer: Berlin, 2005; 205–216.
4. Christin N, Weigand AS, Chuang J. Content availability, pollution and poisoning. *Proceedings of the ACM E-Commerce Conference*. ACM Press: New York, 2005.
5. ABC. <http://sf.net/projects/pingpong-abc> [25 January 2007].
6. Pennisi E. How did cooperative behavior evolve? *Science* 2005; **309**(5731):93.
7. Jelasiy M, van Steen M. Large-scale newscast computing on the Internet. *Technical Report IR-503*, Vrije University, Amsterdam, 2002.
8. Merkle RC. A digital signature based on a conventional encryption function. *Proceedings of the CRYPTO'87 (Lecture Notes in Computer Science, vol. 3279)*. Springer: Berlin, 1988; 369–378.
9. Breese JS, Heckerman D, Kadie C. Empirical analysis of predictive algorithms for collaborative filtering. *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann: San Francisco, CA, 1998.
10. Broder A, Mitzenmacher M. Network applications of bloom filters: A survey. *Proceedings of the 40th Conference on Communication, Control, and Computing*. University of Illinois: Urbana-Champaign, IL, 2002.
11. Roozenburg J. Secure decentralized swarm discovery in TRIBLER. *Master's Thesis*, Delft University of Technology, 2006.
12. Plissonneau L, Costeux J-L, Brown P. Analysis of peer-to-peer traffic on ADSL. *Passive and Active Measurements Conference (Lecture Notes in Computer Science, vol. 3431)*, Dovrolis C (ed.). Springer: Berlin, 2005; 69–82.
13. Voulgaris S, van Steen M. Epidemic-style management of semantic overlays for content-based searching. *Proceedings of Euro-Par 2005 (Lecture Notes in Computer Science, vol. 3638)*. Springer: Berlin, 2005; 1143–1152.
14. Wong J. Enhancing collaborative content delivery with helpers. *Master's Thesis*, The University of British Columbia, 2004.
15. Garbacki P, Iosup A, Epema DHJ, van Steen M. 2Fast: Collaborative downloads in P2P networks. *Proceedings of the 6th IEEE International Conference on Peer-to-Peer Computing*. IEEE Computer Society Press: Los Alamitos, CA, 2006; 23–30.
16. Fast A, Jensen D, Levine BN. Creating social networks to improve peer-to-peer networking. *Proceedings of the 11th ACM SIGKDD*. ACM Press: New York, 2005.
17. Borch N. Social Peer-to-Peer for social people. *Proceedings of the International Conference on Internet Technologies and Applications*. North East Wales Institute: Wrexham, 2005.
18. BTslave project. <http://btslave.sourceforge.net> [25 January 2007].