A Reputation based System to Overcome Malicious Behavior in Peer-to-Peer Networks

A.M. Anisul Huq
Faculty Member, Department of CS,
American International University - Bangladesh (AIUB),
Dhaka - 1213, Bangladesh

Mosharraf Hossain Khan
Daffodil International University,
Bangladesh

ABSTRACT

In recent times, the growth in the number of subscribers of peer to peer networks has been phenomenal. Anonymity being a character of such networks also gave rise to the number of free-riders and malicious behaviors. Though free riders consume network bandwidth and decrease the network performance by displaying selfish behavior, they are not a serious threat for the rest of the co-operative peers. Malicious peers on the other hand, spread viruses, worms, Trojans in the network, provide misleading feedback and try to disrupt the existing trust among the peers. Therefore, it is absolutely essential that a peer has reliable reputation information about other peers in order make informed decisions (e.g., who to download files from, who to serve content, which is a malicious node etc.). In this paper, we have proposed a reputation system that uses objective criteria to track each peer's contribution in the system and allows peers to store their reputations locally. In our opinion, such measures will eventually root out malicious peers from the system.
[1]

Keywords

reputation, peer-to-peer network, trust, eclipse attack

1. INTRODUCTION

The earliest application of peer-to-peer (P2P) was for newsgroups (USENET) and to exchange messages (FidoNet) [2]. Then Napster emerged. With its free music sharing platform and subsequent battle with the big music corporations brought the whole concept of P2P networks into limelight.

P2P networks are primarily used for sharing files and more recently for distributed computations. But studies have shown that the majority of file sharing users do not offer any files for upload, but only download from others [4, 3, 18, 20, 21]. Those who do share are doing it mostly out of ignorance, for not even being aware of it. Or maybe they are indifferent about it, as their uplink bandwidth would simply go unused otherwise and their own download service quality does not suffer from uploads [3]. The presence of malicious peers is further complicating matters and this is the main concern of this paper. They pose a bigger threat because their main goal is to destroy data [25] or damage the infrastructure by propagating worms in the system [26]. As all the systems in a p2p network run the same software, it is very easy for an attacker to compromise the whole network by finding a single exploitable security hole in that software [17].

In order to address these two challenges, we have proposed a reputation mechanism. However, there is no universal agreement on the definition of reputation. In this paper, we have adopted the following working definition:

Reputation: a peer's belief in another peer's capabilities, honesty and reliability *based on recommendations* received from other peers. This *recommendation facility* is also extended to individual files. Reputation can be centralized, computed by a trusted third party, like a Better Business Bureau; or in our case, it is decentralized, computed independently by each peer after asking other peers for recommendations.

The structure of the paper is as follows: background knowledge is presented in Section 2. On the subsequent section, specific p2p attacks and respective defense mechanisms are described. In section 4, we discuss our proposed scheme of reputation. The fifth section analyzes the result and provides relevant comparison.Finally, section 6 concludes the paper.

2. BACKGROUND

Peer-to-peer (P2P) networks have introduced a new paradigm in content distribution. Each peer is both a client and a server in these networks. Users are drawn to these networks due to the ability to locate a wide variety of multimedia content. Currently, there are several different architectures for P2P networks:

(1) ***Centralized:*** There is a constantly-updated directory hosted at central locations. Nodes issue queries to this central directory server to locate which other nodes hold the desired files. Such centralized approaches do not scale well and have single points of failure.
(2) ***Decentralized but Structured:*** Such systems do not have central directory server, but possess significant amount of structure by the way of P2P overlay topology. Such a topology is tightly controlled and files are placed at specific locations that enables

[1] This paper is an extension of the work produced by the first author as a part of his Master's course work done at Aalto University, Helsinki, Finland. It is available at: `http://www.cse.hut.fi/en/publications/B/5/papers/huq_final.pdf`. Note that, this work is not a peer reviewed publication.

queries to be satisfied easily. Such structured P2P systems may use Distributed Hash Table (DHT) as a base, in which data object (or value) location information is placed deterministically. At the peers are identifiers that correspond to the data object's unique key.

(3) ***Decentralized and Unstructured:*** These are systems in which there is neither a central directory nor there is any control over the network topology or file placement. In such a system, nodes join the network by following some loose rules. Here, content retrieval involves a content search and a content download phase. To search for the desired content, a peer generates a query with appropriate keywords and sends it to all the peers that it is directly connected to. The peers who process this queried file only respond if they have the content. These peers in turn will forward this request to those peers with which they are directly connected to. However, this forwarding will depends on the time-to-live (TTL) of the query. The forwarding will continue until the TTL is exhausted. Once the querying peer receives all the replies, it selects a peer to download the content from using either HTTP or a TCP connection [16, 11].

In this paper, we have built a reputation system for decentralized, unstructured P2P systems. We do so because **(1)** these systems are used by large communities of Internet users and **(2)** these systems have not yet been subject to much serious research, except for empirical studies [16].

During both the content search and content download phase, ample cooperation among peers is necessary. The success of the search phase depends on whether the other peers are online or not, if they agree to search for the content within their shared directories, and also forward the query further depending on its hop count. The success of the download phase requires that the chosen peer be online and serve the content when requested. Therefore it is essential for a peer to know the reputations of other peers in the system. Otherwise malicious activity will become rampant and virus, malware and fake files will spread like wild fire within the system. Over the years, many reputation systems have been tested and developed; some of which we will discuss in brief now.

Before it became inactive in 2012, *Kazaa* defined a participation level for each peer based on the MBytes it transferred and the integrity of the files it served. Downloaded file's integrity was labeled as either excellent, average, poor, or delete by the user. The peers are assigned in low, medium, and high category; based on the ratio of Mbytes uploaded and downloaded and the integrity rating of the files. User participation level varied between 0 and 1000. A new user starts at a medium participation level of 100. This participation score was utilized in prioritizing among peers during high demand periods. [11]

Aberer and Despotovic's [1] binary trust model labels a peer as either trustworthy or not trustworthy. It assumes maliciousness is an exceptional occurrence and the peers only store information about their view of the malicious behavior of the peers they interact with. The overall trust is computed on the fly by querying appropriate peers. This system does not have any defense against inserting fraudulent complaints about peers.

Demiani et. al.'s [8] proposal kept separate local repositories for resources and peers. Peers updated their local repositories for the resources and their offer upon finishing transactions. The criteria for such updates are subjective. To compute trust values for resources and the peers on the fly using votes, they enhance the 2 phase search and download protocol into a 5 phase protocol that they developed called, XRep.

NICE [14] is a platform that gains access to the remote contents by bartering local resources. The reputation in NICE is stored as a cookie which can take any real values in the [0; 1] interval and is based on a peer's subjective satisfaction from its transaction. As the peers store their own reputations, cooperation from other peers is only needed when reputation is computed.

PeerTrust [25] is also a feedback based trust management system where reputation is computed based on three factors: 1) the amount of satisfaction received by the other peers in the system, 2) the total number of interactions, and 3) a balancing factor to offset the impact of malicious peers that misreport other peers' service. Each peer is mapped to maintain a small database that stores a portion of the global trust data. Maliciousness is countered by having multiple peers responsible for storing the same database. Voting can be used if these databases differ. Trust is computed on the fly through querying potentially multiple databases.

Before we discuss our proposed reputation system, we need to elaborate on the nature of the threats posed by malicious peers and how reputation based mechanisms can help alleviate them. In the next section, we will discuss several malicious p2p attacks and their defenses in the context of reputation.

3. SPECIFIC P2P ATTACKS AND DEFENSES

Basically, there have been two broad categories of attacks on the P2P networks. In the first type, attackers target the data circulating in the P2P networks, e. g. by corrupting it or making it unavailable for other peers. In the other type, attack involves making the network as slower or inefficient as possible. This sort of attack is generally done by exploiting the under lying weakness of the routing protocol. Depending on the attacker's objective, he may choose to attack from any one direction or from both [17].

Now, in many cases attacks of one type can trigger the other. For example, by corrupting files an attacker can prompt users to download more copies of a much sought after file, thus slowing down the network. The opposite is also true. In case of eclipse attacks networks are blocked (hence inefficient) making data inaccessible which is an objective of the first type of attack [17].

The possibility of attack is enormous in P2P networks. We now give an analysis of the most common type of attacks along with the traditional defense mechanisms that are currently employed against them.

3.1 Rational Attacks

By the term "rational" we indicate to those peers who will attempt to maximize their consumption of system resources (one may choose to call them "selfish") while minimizing the use of their own. Research shows that a big portion of the peers are of this type [25]. Peers with limited bandwidth capabilities are more prone to this tendency. Also in sharing copy right material a peer might find itself in legal problems [12]. These are good enough reasons to motivate nodes in becoming "self-interested". If a large number of nodes behave in this way, it will cause the overall performance of the network to plummet.

3.1.1 Defenses. With perfect global knowledge of every peer's reputation, a node would receive incentive for cooperation. Any time it cheats, information would be immediately available to all of its peers. EigenTrust describes how such a global systems can be built [13].

3.2 File Corruption

As the name suggests this is an attack against data in the P2P network. The objective here is to replace a file in the network with a false one. In order to attack in this manner, malicious nodes will falsely claim of owning a file, and upon a request will respond with a corrupt file. Moreover, all messages passing through malicious peer can be corrupted (similar to a man-in-the-middle attack) giving these files a high availability [17]. Surprisingly, it is not only individuals or a rouge group of peers that are involved in file corruption attacks. It has also been reported that, the music industry has massively dumped corrupt and fake contents into the P2P networks [7, 15, 17].

3.2.1 Defenses. Though corruption attacks sound pretty dangerous, Dumitriu et al. [10] argue they do not pose a serious threat to the P2P networks. The main problem is that P2P applications often run in the background. When a polluted file is downloaded, it stays available for a while before it is checked by the user and discarded. Our proposed reputation system will accelerate this process, as it also puts scores to individual files. After a period of time, all polluted files will be removed and the authentic files will become more available then the corrupted ones.

3.3 Sybil Attack

Sybil is of the second type of attack that we mentioned at the beginning of this section. It is about making the network cripple and inefficient. Generally, in a structured P2P network, user identifiers (IDs) uniquely identify participant endpoints (nodes). Such structure reduces search times by mapping content directly onto nodes based on IDs. For this reason, the assignment and use of IDs is essential to correct operation of the network [19]. Now, it is very much possible that a single malicious peer can generate multiple shadow identities and thus gain control over a part of the network [9]. Once this has been accomplished, the attacker can gain access to certain files and may decide to corrupt those. If the attacker can position his false identities in a strategic way, the damage can be considerable. He might choose to continue to an eclipse attack, or slow down the network by re-routing all queries in a wrong direction.

3.3.1 Defenses. Douceur et al. [9] have shown that, with a central trusted authority P2P systems can defend against Sybil attacks. In addition to a centrally trusted authority, several papers have proposed a complicated public-private key based protocol [22] where each peer must sign its messages, and respond to a challenge by the authority at random. It is clear that an attacker simulating many identities would need enormous resources in order to be able to answer all the challenges periodically submitted to each of his identities. While this certainly tries to solve the problem, it is unsatisfactory. It breaks the P2P model by reintroducing a centralized point of failure, which can easily be attacked.

Cheng and Friedman [6] has evaluated the vulnerability of reputation systems to the Sybil attack and has classified these reputation systems as either symmetric or asymmetric. [2] In an asymmetric system, there are specifically trusted nodes from which all reputation values propagate. Alternatively, each entity separately computes a trust value along their unique paths to every other identity in the system. Since the trusted nodes cannot be impersonated, no Sybil attacker can create a duplicate graph as explained above in the symmetric case. This trust value can change over time as the entity interacts with and observes the behavior of different identities. Our proposed reputation system is quite similar to this approach.

3.4 Eclipse Attack

In an overlay network, each node maintains links to a relatively small set of peers called neighbors. All communication within the overlay (it may be related to maintaining the overlay or to application processing) occurs on these links [23]. The overlay network's integrity depends on the ability of correct nodes to communicate with each other over a sequence of overlay links. In an Eclipse attack [5, 23, 24] a modest number of malicious nodes conspire to fool correct nodes into adopting the malicious nodes as their peers, with the goal of positioning themselves along strategic routing paths of the P2P network. Once an attacker has done this, he can separate the network in more than one sub networks. After that, if a peer wants to communicate with a peer from some other sub network, its message must at a certain point be routed through one of the attacker's nodes. The attacker thus "eclipses" each sub network from the others' view [23]. The following figure gives a clear idea of what happens.

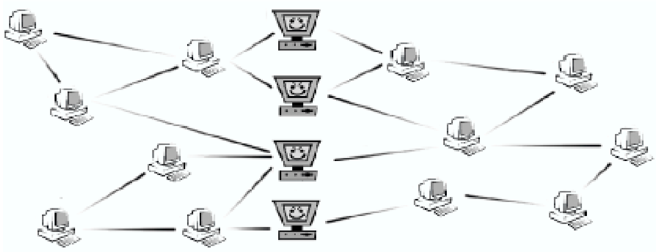

Fig. 1. An Eclipse attack: the malicious nodes have separated the network in 2 sub-networks. [17]

If the P2P network is based on a decentralized overlay network then nodes will periodically discover new neighbors by consulting the neighbor sets of existing neighbors. Malicious nodes can exploit this by advertising neighbor sets that consist of only other malicious nodes. Thus, a small number of malicious nodes with legitimate identities are sufficient to carry out an Eclipse attack. Castro et al. identify the Eclipse attack as a threat in structured overlay networks [5].

3.4.1 Defenses. As we have said earlier, Eclipse attack is closely related to the Sybil attack [9, 23] and a successful Sybil attack can lead to an Eclipse attack. Therefore according to *Cheng and Friedman* [6], if we are successful in defending against a Sybil attack (using an asymmetric reputation system) then it might be possible to stop an impending Eclipse attack.

The method introduced in [22] can also be used to prevent eclipse attack. According to this method, a node that mounts an Eclipse attack must have a higher than average node degree. *Singh et al.* [22] argues that enforcing a node degree limit by auditing is an effective defense against Eclipse attacks.

[2] *Cheng and Friedman* [6] has proven that symmetric reputation systems are susceptible to Sybil attacks and is therefore irrelevant in our current discussion.

From the above discussion we can conclude that, if we can track the reputation of each node, as well as individual file; it will be easier to differentiate between good/cooperative and malicious nodes. P2P network users will only interact with only those peers with high reputation scores; while avoiding the less reputed ones. It will substantially decrease the presence of malicious nodes in the network.

4. PROPOSED REPUTATION SYSTEM

In [12] we have seen that maximum of reputation based system for P2P network works in centralized manner; whereas maximum of the P2P network are decentralized in nature. Therefore, we are proposing a credit-debit based reputation suitable for decentralized P2P network.

4.1 credit-debit system:

Our proposed reputation system is a credit and debit based system which we will now discuss. In this paper, we have coined the terms, SanitizedFile and UnsanitizedFile for our convenience. The first one represents an un-corrupted, genuine file; while second one indicates to fake/ corrupted / malicious file.

4.1.1 SanitizedFile Credit. For every SanitizedFile, the node will get credit point. This credit point will be added to calculate the overall reputation score. This credit point will be given by the users who downloaded content from this node. If the user does not give any feedback, zero (0) will be added to the node. We are aware that many of the users are reluctant to give any feedback that is why we are also proposing to add reputations component for users also. This reputation component tracks the reputation of the users who download files from other nodes. This component will add one if the user gives feedback and give minus one if the user does not give feedback. This reputation score of the user will influence the download speed for each user. Thus user will be motivated to give feedback.

4.1.2 UnsanitizedFile Debit. For every UnsanitizedFile the provider node will get debit point. It will be subtracted from the overall score.

The reputation score will be file based. By this we mean, the reputation will be calculated for each file and will eventually be summed up to calculate the overall reputation. The scores will range from 1 to 5. For any positive peer review the score will increase by the same amount of rating and for any negative review the given rating will be first multiplied by 2 then subtracted from the overall score. In our simulation we have observed that, this policy of putting more weight on negative review provides better results than putting same weight for credit and debit.

This strategy has to be each file based so that one SanitizedFile with high download cannot hide multiple UnsanitizedFiles in the system. For example, if a SanitizedFile has 50 downloads with an average of 3.5 rating and 5 UnsanitizedFiles, each with 10 downloads with -1 rating then the system will still have good reputation (overall reputation =((50 x 3.5) (5 x 10 x 2))). But the real situation is that the system is disseminating 5 malicious files and only 1 SanitizedFile.

The reputation score is shown in two tables as one node can both download and upload content. In table 1, we keep track of the overall score for download reputation (DR) and upload reputation (UR). In table 2, we keep track of each file. It will show the average score for each file, which is called File Reputation (FR).

Table 1: Overall Reputation Score

Name	Name
Download Reputation	5
Upload Reputation	3.5

Table 2: Reputation Score for Each Individual File

Sr. No	Name of the File	Reputation Score
1.	Photo.jpg	2.3
2.	Business plan.doc	4.5
3.	Movie.mpeg	4
4	Virus.exe	-3
. . .		
100.	Lecture.pdf	3.3

4.1.3 Download Reputation (DR). If the user provides necessary feedback, then he/she will get 1 point. If the user refrains then his/her point gets deducted. Therefore, DR score will be the summation of all the scores.

$$DR = \sum(Given\ feedback) - \sum(not\ given\ feedback)$$

4.1.4 Upload Reputation (UR). Upload reputation is calculated as the average reputation of latest 100 downloaded files.

$$UR = \frac{\sum(File\ Reputation\ (FR))}{N}$$

Here, *N* is the number downloaded files.

4.1.5 File Reputation (FR). File reputation is calculated as the summation of all the feedback score for each file.

$$FR = \sum(Positive\ Score) - 2\ *\ \sum(Negative\ Score)$$

Security of the reputation system is of great concern today. Major threats faced by this kind of reputation systems include authentication, trust and non-repudiation. In order to handle such security threats, our envisioned system should use a public-private key based infrastructure. The counter that keeps the most recent reputation score for each peer is updated and stored in the enrolled peer's local software. The local storage allows for fast retrieval of reputations. Each peer interested for enrollment in the reputation computations generates a (public, private) key pair and registers it with any publicly available public key system (e.g, PGP system). The digest of the public key is used to identify the peer. Thus we can identify the fake nodes in the system.

However, this security mechanism has its downsides. For one, some peers may not want to get their reputation tracked for privacy reasons. Existing designs of P2P networks do not provide peer anonymity and our goal in this paper is not to propose alternate designs of P2P networks. As a result, the reputation tracking presented here does not address anonymity issues in such tracking. Also, the reputation system involves additional overheads to keep the most up-to-date view of each peer's reputation which some peers may not want to incur. For these reasons, enrollment in the reputation computations is voluntary. Peers who choose not to enroll always maintain a default reputation score of 0.

5. RESULT AND COMPARISON

We have simulated our proposal using Arena. Number of nodes in the simulation is 50 and 15 % of the node are malicious. We have run the simulation upto 5000 downloads. The downloading procedure is random. Any node can download files from its neighbor. The choice of node is random. As we have proposed to double the negative effect, we wanted to see the effect. In figure 2, we see that doubling the negative rating decreases number of malicious files.

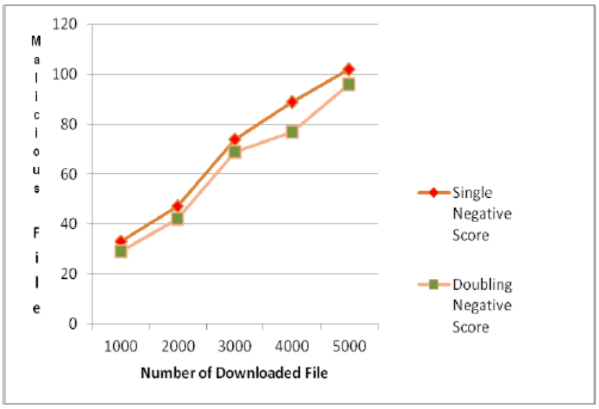

Fig. 2. Single Negative Vs Double Negative Score.

Another important feature of our proposed reputation system is the keeping reputation score for each files. In our simulation it shows that keeping reputation score for each file decreases the download of malicious files. Figure 3, shows that significant improvement can be achieved by keeping record of each files separately.

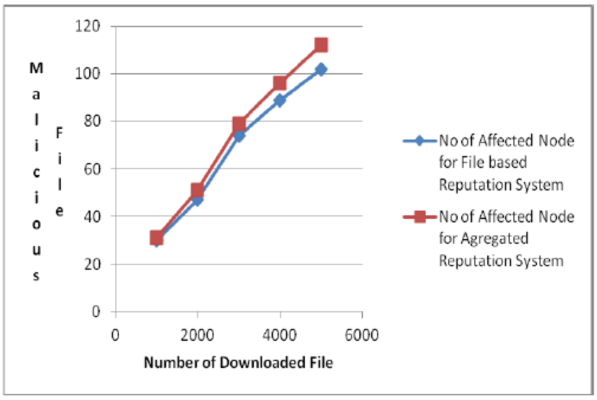

Fig. 3. Single File based Vs Aggregated Reputation

6. CONCLUSION

This paper proposes a reputation system for decentralized unstructured P2P networks. The simulation results show that the proposed system provides better result than the approaches used in other similar systems. Our reputation scores are stored and maintained in the local node, negating the need for any central server or infrastructure. Through our simulated comparison we have shown that, doubling the negative rating and keeping reputation scores for each individual file decreases malicious activity in P2P networks.

7. REFERENCES

[1] Karl Aberer and Zoran Despotovic. Managing trust in a peer-2-peer information system. In *Proceedings of the tenth international conference on Information and knowledge management*, pages 310–317. ACM, 2001.

[2] Abiola Abimbola, Qi Shi, and Madjid Merabti. Using intrusion detection to detect malicious peer-to-peer network traffic. *Proc. of PGNet*, 2003.

[3] Torsten Ackemann, Cecilia Mascolo, and Wolfgang Emmerich. Lightweight incentives for peer-to-peer networks. 2003.

[4] Eytan Adar and Bernardo A Huberman. Free riding on gnutella. *First Monday*, 5(10), 2000.

[5] Miguel Castro, Peter Druschel, Ayalvadi Ganesh, Antony Rowstron, and Dan S Wallach. Secure routing for structured peer-to-peer overlay networks. *ACM SIGOPS Operating Systems Review*, 36(SI):299–314, 2002.

[6] Alice Cheng and Eric Friedman. Sybilproof reputation mechanisms. In *Proceedings of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, pages 128–132. ACM, 2005.

[7] Nicolas Christin, Andreas S Weigend, and John Chuang. Content availability, pollution and poisoning in file sharing peer-to-peer networks. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 68–77. ACM, 2005.

[8] Ernesto Damiani, De Capitani di Vimercati, Stefano Paraboschi, Pierangela Samarati, and Fabio Violante. A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *Proceedings of the 9th ACM conference on Computer and communications security*, pages 207–216. ACM, 2002.

[9] John R Douceur. The sybil attack. In *Peer-to-peer Systems*, pages 251–260. Springer, 2002.

[10] Dan Dumitriu, E Knightly, Aleksandar Kuzmanovic, Ion Stoica, and Willy Zwaenepoel. Denial-of-service resilience in peer-to-peer file sharing systems. In *ACM SIGMETRICS Performance Evaluation Review*, volume 33, pages 38–49. ACM, 2005.

[11] Minaxi Gupta, Paul Judge, and Mostafa Ammar. A reputation system for peer-to-peer networks. In *Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video*, pages 144–152. ACM, 2003.

[12] Bill Horne, Benny Pinkas, and Tomas Sander. Escrow services and incentives in peer-to-peer networks. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 85–94. ACM, 2001.

[13] Sepandar D Kamvar, Mario T Schlosser, and Hector Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the 12th international conference on World Wide Web*, pages 640–651. ACM, 2003.

[14] Sang-Rim Lee, Rob Sherwood, and Bobby Bhattacharjee. Cooperative peer groups in nice. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and*

Communications. IEEE Societies, volume 2, pages 1272–1282. IEEE, 2003.

[15] Jian Liang, Rakesh Kumar, Yongjian Xi, and Keith W Ross. Pollution in p2p file sharing systems. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, volume 2, pages 1174–1185. IEEE, 2005.

[16] Qin Lv, Pei Cao, Edith Cohen, Kai Li, and Scott Shenker. Search and replication in unstructured peer-to-peer networks. In *Proceedings of the 16th international conference on Supercomputing*, pages 84–95. ACM, 2002.

[17] Baptiste Pretre. Attacks on peer-to-peer networks. *Dept. of Computer Science Swiss Federal Institute of Technology (ETH) Zurich Autumn*, 2005.

[18] Matei Ripeanu, Ian Foster, and Adriana Iamnitchi. Mapping the gnutella network: Properties of large-scale peer-to-peer systems and implications for system design. *arXiv preprint cs/0209028*, 2002.

[19] Hosam Rowaihy, William Enck, Patrick McDaniel, and Thomas La Porta. Limiting sybil attacks in structured p2p networks. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 2596–2600. IEEE, 2007.

[20] Stefan Saroiu, P Krishna Gummadi, and Steven D Gribble. Measurement study of peer-to-peer file sharing systems. In *Electronic Imaging 2002*, pages 156–170. International Society for Optics and Photonics, 2001.

[21] Subhabrata Sen and Jia Wang. Analyzing peer-to-peer traffic over large networks. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, volume 257, page 258, 2002.

[22] Atul Singh, Miguel Castro, Peter Druschel, and Antony Rowstron. Defending against eclipse attacks on overlay networks. In *Proceedings of the 11th workshop on ACM SIGOPS European workshop*, page 21. ACM, 2004.

[23] Atul Singh et al. Eclipse attacks on overlay networks: Threats and defenses. In *In IEEE INFOCOM*. Citeseer, 2006.

[24] Emil Sit and Robert Morris. Security considerations for peer-to-peer distributed hash tables. In *Peer-to-Peer Systems*, pages 261–269. Springer, 2002.

[25] Li Xiong and Ling Liu. Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities. *Knowledge and Data Engineering, IEEE Transactions on*, 16(7):843–857, 2004.

[26] Runfang Zhou, Kai Hwang, and Min Cai. Gossiptrust for fast reputation aggregation in peer-to-peer networks. *Knowledge and Data Engineering, IEEE Transactions on*, 20(9):1282–1295, 2008.