The P2P War: Someone Is Monitoring Your Activities!

Anirban Banerjee, Michalis Faloutsos, and Laxmi Bhuyan

Department of Computer Science and Engineering
University of California, Riverside
Riverside, CA 92521
{anirban,michalis,bhuyan}@cs.ucr.edu

Abstract. In an effort to prosecute P2P users, RIAA and MPAA have reportedly started to create decoy users: they participate in P2P networks in order to identify illegal sharing of content. This has reportedly scared some users who are afraid of being caught. The question we attempt to answer is how prevalent is this phenomenon: how likely is it that a user will run into such a “fake user” and thus run the risk of a lawsuit? The first challenge is identifying these “fake users”. We collect this information from a number of free open source software projects which are trying to identify such IP address ranges by forming the so-called blocklists. The second challenge is running a large scale experiment in order to obtain reliable and diverse statistics. Using Planetlab, we conduct active measurements, spanning a period of 90 days, from January to March 2006, spread over 3 continents. Analyzing over a 100 GB of TCP header data, we quantify the probability of a P2P user of being contacted by such entities. We observe that 100% of our nodes run into entities in these lists. In fact, 12 to 17% of all distinct IPs contacted by any node were listed on blocklists. Interestingly, a little caution can have significant effect: the top five most prevalent blocklisted IP ranges contribute to nearly 94% of all blocklisted IPs and avoiding these can reduce the probability of encountering blocklisted IPs to about 1%. In addition, we examine other factors that affect the probability of encountering blocklisted IPs, such as the geographical location of the users. Finally, we find another surprising result: less than 0.5% of all unique blocklisted IPs contacted are owned explicitly by media companies.

Keywords: Peer-to-peer networks, Gnutella, RIAA, User monitoring.

1 Introduction

Organizations like the RIAA and MPAA, representing content providers, have escalated their fight against illegal P2P content sharing [2], [13], [14], [15], [21], [22] with the use of fear: there have been a number of lawsuits against individual P2P users [3], [4], [5], [6]. For greater effect, these organizations and their collaborators have also started “trawling” in P2P networks: creating “fake users” which

participate in the network and thus identify users who contribute towards illegal content sharing. However, the extent of this deployment tactic has not been quantified up to now, and this forms the focus of our work.

In response to this approach, the P2P community has spawned several projects which attempt to: (a) identify such “fake users”, and (b) enable P2P users to avoid them. In more detail, there is a community based effort to maintain lists of suspect IP address ranges, which are called *blocklists*. Blocklists are published by organizations which provide anti-RIAA software or by groups which focus on security [9]. Additionally, a number of free, open-source, software projects enable P2P users to avoid these blocklisted IPs automatically and are integrated with the most popular P2P clients using BitTorrent, eDonkey/ eMule, Gnutella networks [1], [8], [9], [30], [17], [26]. Note that it is not our intention here to examine how accurate and comprehensive these lists are, though this would be interesting and challenging future work. What we claim is that, *the information we use in our research, is readily available to P2P users and is used by them.* [1].

The question we attempt to answer is, how prevalent is the phenomenon of fake users. Simply put, how likely is it that a user will run into such a “fake user” without using blocklists? The answer to this question can lead us to: (a) understand the effort that content providers are putting in trawling P2P networks, and (b) justify the effort of the P2P community to isolate “fake users”. Hereonwards, we refer to IP ranges of fake users listed on these blocklists as blocklisted IPs and users exchanging data with them as **being monitored**. The intention of blocklists is to identify such “monitoring” entities, however all IP ranges listed on blocklists are not monitoring users, but we assume the “worst” case scenario. We say that a user **hits** blocklisted IPs every time a user receives or sends a piece of data (part of a file) to that IP range. Organizations employing these blocklisted IPs are referred to as blocklisted entities. To the best of our knowledge, such measurements have not been collected before.

We conduct what seems to be the first study on the probability with which P2P users are being monitored. We employ PlanetLab [12] for a period of 90 days and customize a Gnutella client (mutella version 0.4.5) to automatically initiate meaningful queries and collect statistics from the Gnutella network. Each client initiates 100 queries for popular song found in prominent music charts [36], [29], [28]. We collect and analyze nearly 100GB of TCP header data. We then examine the observed IP addresses using the most popular blocklists on the Internet [1], [9], [30].

Our results can be summarized as follows:

1. **Consequence of ignoring blocklists:** A user without any knowledge of blocklists, will almost certainly be monitored by blocklisted IPs. We found that **all** our clients exchanged data with blocklisted IPs. In fact, of all distinct IPs contacted by any client, 12-17% were found to be listed on blocklists.
2. **A little information goes a long way:** We find that avoiding just the top 5 blocklisted IPs reduces the chance of being monitored to about 1%. This is a consequence of a skewed preference distribution: we find that the top 5

blocklist ranges encountered during our experiments contribute to nearly 94% of all blocklist hits.

3. **Most blocklisted IPs belong to government or corporate organizations:** We quantify the percentage of hits to blocklisted entries of each type, i.e. government and corporate, educational, spyware proliferators and Internet advertisement firms. We find that the number of hits which belong to government and corporate lists, is approximately 71% of total number of hits, nearly 2.5 times more than educational, spyware and adware lists put together. Interestingly, some blocklists mention unallocated IP ranges called BOGONS, which we discuss later.
4. **Very few blocklisted IPs belong directly to content providers:** We find that 0.5% of all blocklisted IPs hits could actually be traced back to media companies, such as Time Warner Inc. However, it is an open question whether other blocklisted IPs are indirectly related to content providers.
5. **Geographical bias:** We find that there is geographical bias associated with how users hit entities listed on blocklists. The way in which users located on the two opposite coasts, east and west, of mainland US, Europe and Asia, hit blocklisted entities is quite different.
6. **Equal opportunity trawling:** We find that Ultra-peers (UPs)¹ and leaf nodes have equal probability of associating with a blocklisted IP, with less than 5% variation in the average number of distinct blocklisted IPs. This comes in contrast to the popular belief that UPs are monitored more aggressively by blocklisted entities [10], [11], than leaf users.

The rest of the paper is organized as follows. Section II presents relevant literature, followed by Section III which discusses the experimental setup and blocklisted entries. Section IV investigates geographical bias and section V addresses the Ultra/Super peer versus leaf node debate.

2 Relevant Literature

A plethora of P2P networks, such as FastTrack, Gnutella [14], BitTorrent, eMule/Donkey and many others are prevalent in the Internet. Freely available P2P clients for nearly all operating systems generate significant amounts of traffic criss-crossing the Internet [13], [15]. These networks have recently been touted as the future for content distribution technologies [16], and for similar exciting and promising applications. However, these overlay networks act as significant enablers in the movement of copyrighted material over the web. Organizations such as the RIAA and MPAA have been vociferous in their support for anti-P2P policies since it is the companies represented by these organizations that supposedly lose out on revenue due to the exchange of copyrighted songs and movies [5], [7].

Recently, a slew of reports in the electronic and print media have led to members of P2P communities pondering over the ramifications of such illegal resource

¹ Ultra-peers are high bandwidth nodes that act as local centers, facilitate low bandwidth *leaf* nodes, and enable the scalability of gnutella-like networks.

sharing [18]. To mitigate such a threat of possible lawsuits, users have resorted to downloading and deploying anti RIAA/MPAA software. These programs block computers owned by such organizations from accessing users on the P2P networks [8], [1], thereby effectively alienating them from quorums of P2P users. This prevents them from gaining critical information leading to generation of detailed user behavior log files which may be used for legal action. The number of such free software, easily available from popular websites is large. Many variants exist for different clients, networks and Operating Systems.

Previous work on modeling and analysis of P2P systems [24], [25], have focused on developing a viewpoint based on performance metrics of such overlay systems. Our work differs greatly from these important earlier research efforts. We conduct research to specifically ascertain if the organizations like the RIAA are active on P2P networks or not. We quantify the probability of a P2P user of being monitored by entities listed on the most popular blocklists. Also, we identify if there is any geographical bias associated with observing how P2P users run up against blocklisted entities. To the best of our knowledge, we believe that our research is the first which specifically targets an in-depth study of whether such a threat is a reality for a generic P2P user. Moreover, our work is significant for understanding *who do we talk to* while sharing copyrighted resources on these P2P networks. Additionally, we intend to verify reports suggesting that some so-called organizations enlisted by the RIAA *target UPs in preference to leaf nodes* [10], [11], in order to break the backbone of the entire overlay structure.

3 Who Is Watching?

In this section we discuss the experimental setup we employ followed by a synopsis of our findings regarding which blocklisted entities are most prevalent on P2P networks.

Experimental set-up. We initiate our experiments to emulate a typical user and yet be able to measure large scale network-wide inter-node interaction characteristics of P2P networks. We measure statistics based on trace logs compiled from connections initiated using PlanetLab. The duration of measurements spanned more than 90 days, beginning January 2006. We initiate connections using nodes spread not only across the continental US but also Europe and Asia in order to determine any geographical nuances associated with which blocklisted entities seem to be more active than others, in specific locations. We were able to customize mutella 0.4.5 clients [27], a vanilla console based Gnutella client, and initiate connections to the Gnutella network. Moreover, clients were made to switch interchangeably from UP to leaf modes in order to verify if network wide inter-node behavior of UPs is significantly different from leaf nodes.

Search strings used for probing the P2P network were compiled as a list of popular songs, from Billboards hot 100 hits [28], top European 50 hits [29] and Asian hits [36]. Each node injected about 100 queries during every run. In the process, we analyzed more than 100GB of TCP header traces by using custom scripts and filters to extract relevant information which helps us develop a deeper

insight into who do we interact with while sharing resources on P2P networks. Note that *all files stored as a result of our experiments on PlanetLab nodes, were completely removed and never used*. Similarly no content was downloaded to local UCR machines for storage.

Before we present results obtained from our measurements we must discuss what BOGON IPs [34] mean as they hold special significance to the collected information. BOGON is the name used to describe IP blocks not allocated by IANA and RIRs to ISPs and organizations plus all other IP blocks that are reserved for private or special use by RFCs. As these IP blocks are not allocated or specially reserved, such IP blocks should not be routable and used on the internet, however some of these IP blocks do appear on the net primarily used by those individuals and organizations that are often specifically trying to avoid being identified and are often involved in such activities as DoS attacks, email abuse, hacking and other security problems.

The majority of the most active blocklisted entities encountered are hosted by organizations which want to remain anonymous. Table 1 lists the top fifteen entities we encounter on the P2P network while exchanging resources, throughout the complete duration of our active trace collection. Surprisingly, we find these entities operate from BOGON IP ranges. This observation is made on the basis of the various popular blocklist resources, and suggests that *these sources deliberately wish to conceal their identities while serving files on P2P networks*, by using up IP ranges which cannot be monitored down using an IP-WHOIS lookup to locate the operator employing these anonymous blocks. Only three out of the top fifteen entries in table 1 do not use unallocated BOGON IP blocks and are listed on PG lists [1]. The rest of the BOGON entities are listed on both Trustyfiles [30] and Bluetack [9] lists. Most of the BOGON IP ranges point to either ARIN or RIPE IP ranges. We must however mention that these BOGON IP ranges were found to point back to these generic network address distribution entities at the time of our experiments. It is quite possible that these ranges may have now been allocated to firms or individuals and may no longer remain anonymous.

Content providers part of the RIAA do not participate in large scale eavesdropping into P2P networks using their own IPs. We observe that a whopping 99.5% of blocklisted IPs contacted either belong to BOGON, commercial entities, educational institutions and others. Among all blocklisted IPs contacted, about 0.5% could actually be traced back to record companies, such as Time Warner Inc. This is a clear indication of the miniscule presence of record companies trawling P2P networks in a proactive manner.

According to popular perception in the P2P community, and discussions on blocklist hosting sites, such as Phoenix Labs [35], the entry FUZION COLO [31], [32] in Table 1, is viewed with distrust, and is understood to propagate self installing malware, and in general as an anti P2P entity. Xeex [33], is more of a mystery. It hosts an inconspicuous site which provides absolutely no information as to what the company is really involved in. Going by the discussion groups

Table 1. Listing of top 15 blacklist entities encountered on P2P network

Rank	Top15HitRanges	Type
1	72.48.128.0-72.235.255.255	Bogon
2	87.0.0.0-87.31.255.255	Bogon
3	88.0.0.0-88.191.255.255	Bogon
4	72.35.224.0-72.35.239.255	FuzionColo
5	71.138.0.0-71.207.255.255	Bogon
6	70.229.0.0-70.239.255.255	Bogon
7	70.159.0.0-70.167.255.255	Bogon
8	70.118.192.0-70.127.255.255	Bogon
9	216.152.240.0-216.152.255.255	xeex
10	216.151.128.0-216.151.159.255	xeex
11	70.130.0.0-70.143.255.255	Bogon
12	87.88.0.0-87.127.255.255	Bogon
13	71.66.0.0-71.79.255.255	Bogon
14	87.160.0.0-87.255.255.255	Bogon
15	70.82.0.0-70.83.255.255	Bogon

Table 2. Listing of top 5 educational and commercial entities encountered on P2P networks

Rank	Top5EducationalHitRanges	Top5CommercialHitRanges
1	152.2.0.0-152.2.255.255- Univ. of N. Carolina	72.35.224.0-72.35.239.255-FuzionColo
2	64.247.64.0-64.247.127.255-Ohio University	216.152.240.0-216.152.255.255-XeeX
3	129.93.0.0-129.93.255.255- Univ. of Nebraska	216.151.128.0-216.151.159.255-XeeX
4	128.61.0.0-128.62.255.255-Georgia Tech	38.113.0.0-38.113.255.255-Perf.SystemsInted2k
5	219.242.0.0-219.243.255.255-CERNET	66.172.60.0-66.172.60.255-Netsentryed2kserver

hosted on the PG website, xeex does turn up frequently in blacklist hits for a large number of users. Other individuals or organizations deliberately employing BOGON IPs to participate in the exchange of resources on P2P networks are certainly attempting to cloak themselves, possibly from the RIAA. Another vein of reasoning would suggest that they could be the ones who keep tabs on what users download.

Table 2 displays the top five entities that registered hits on the educational and research institutions list and the government and commercial organizations lists. We observe that FuzionColo and XeeX appear prominently in this categorization along with two other commercial organizations which host servers on ed2k and Gnutella networks. We find that hits to entities listed on commercial and government blocklists are much more frequent than hits on any other different kind of blocklists such as Internet ad companies, educational institutions and others. Even though the number of IPs which belong explicitly to content providers may be small, the fact that IPs listed on commercial and government blocklists are providing content to P2P users is of concern. The scenario wherein commercial organizations are hired by content providers to collect user profile data in these networks cannot be ruled out. Furthermore, the possibility that these commercial organizations such as the ones listed in table 2 are not aware of P2P traffic emanating from their servers and are too lax about security does not seem very plausible since some of these bocklisted entities kept monitoring our clients nearly every time files were exchanged. It is clear that these commercial IP ranges which serve files to P2P users have a very large cache of popular in-demand media and have extremely low downtime, which seems improbable if in fact the machine were turned into a bot. In fact, the number of hits to

Fig. 1. Classification of blocklist hits according to their type. We observe that hits on the commercial and government blocklist is significantly larger than hits on the other blocklists.

commercial and government blocklisted entities is nearly 2.5 times greater than hits to any other kind of blocklisted IP we were monitored by.

4 Probability

In this section it is our intention to estimate the probability of a typical user of being monitored by entities listed on these blocklists while surfing P2P networks. This gives an idea of how aggressive these lists are and what percentage of entities we talk to while surfing P2P networks are not considered trustworthy. We observe throughout the complete duration of our measurements, **100% of all our nodes were monitored by entities on blocklists and on average 12-17% of all distinct IPs contacted by any of our clients were listed on blocklists.** As illustrated in Fig. 2, the percentage of IPs listed on blocklists which a node is monitored by is quite significant, about 12-17% of all distinct IPs contacted, per node. In fact this trend was reflected throughout the complete duration of measurements, which suggests that the presence of blocklisted entities on P2P networks is not an ephemeral phenomenon.

Popularity of blocklisted IPs monitoring P2P users follows a skewed distribution. We observe this behavior as displayed in Fig. 3a. A small number of entities register a large number of hits while most blocklisted entities are infrequently visible on P2P networks. This fact is of great consequence to users who wish to avoid contact with blocklisted entities and thus reduce their chances of running into anti-P2P entities. *Simply filtering out the five most popular entities on these networks leads to a drastic reduction in the number of hits to them, to the tune of 94%.* This interesting statistic is displayed in Fig. 3b. In fact **avoiding just these top 5 popular IP ranges can reduce the chances of a user being monitored significantly, down to nearly 1%.** Users may use this fact to tweak their IP filters to increase their chances of safely surfing P2P networks and bypassing the most prevalent blocklisted entities. In contrast, a naive user without any information of blocklists will almost certainly be monitored by blocklisted entities. Also, the fact that 100% of all nodes regardless of geographical location were monitored by blocklisted IPs, indirectly points to the completeness of the blocklists we compiled from the most popular sources.

Fig. 2. Percentage of distinct blocklist IPs contacted, per user, out of the total number of distinct IPs logged

Fig. 3. (a) Frequency of popularity of blocklisted IPs, following a skewed distribution. (b) Percentage contribution by Blocklisted IPs. The 5 most popular blocklisted IPs contribute to nearly 94.2% of all blocklist hits.

5 Geographical Distribution

In this section we focus attention towards whether geographical bias if any is observed with respect to blocklisted IPs monitoring our clients from different locations. To achieve this we needed to develop a mechanism allowing us multiple points of entry, geographically speaking, into a P2P network. We employed over 50 different nodes on PlanetLab, encompassing the continental US, Europe and Asia to measure this metric. We monitor individually, PlanetLab nodes located in the continental US and classify nodes situated on the east coast as US-EC and on the west coast as US-WC. This was done to observe if there is any variation in monitoring behavior within mainland US. Surprisingly, we find that measurements gathered from PlanetLab nodes located on US-EC and US-WC do not concur in unison regarding various metrics discussed in the following sections.

Geographical location influences observed monitoring activity. To provide an idea of how blocklisted IPs monitor P2P users over a complete geographical spectrum we present Fig. 4a. We observe that the percentage of blocklisted IP hits is highest in US-WC followed by US-EC, Asia and Europe. *The percentage of hits to blocklisted IPs per node, compared to total hits to IPs contacted by each node, located on the US-WC seems to be nearly twice that of nodes located on US-EC.* Quite obviously, this suggests that users accessing the P2P network from these two vantage points, within the mainland US, encounter different levels

Fig. 4. UP Vs Leaf: (a) Distribution of Blocklisted IPs contacted in different geographical zones. (b) Distribution of Blocklisted IP hits, to Educational lists, in different geographical zones. (c) Distribution of Blocklisted IP hits, to Government and Commercial lists, in different geographical zones.

of monitoring activity. We believe this observed inequality springs from the following reason, that difference in user behavior and possible difference in levels of monitoring activities by entities on the blocklists could directly be responsible for such a skewed trend. Fig. 4b depicts the distribution of blocklisted IP hits from the “educational” range, comprising of academic and research institutions. Again, we observe a similar trend. Nodes located on US-WC notch up a higher percentage of blocklist hits compared to nodes located on US-EC, Asia and Europe. In fact, the difference in measurements between US-WC and US-EC is more than five times than that of readings gathered from US-EC. Fig. 4c depicts the distribution of blocklisted IP hits in the government and commercial domain. Once again, we observe that figures collected for nodes situated on US-WC are higher than nodes on US-EC, Asia and Europe. Given that the period of observation, the UTC time when data was logged, the number of queries input into the P2P network, the order in which queries were injected were identical, we surmise that, throughout the duration of our experiments *the consistent skewed distribution between US-WC and US-EC can be due to difference in user behavior and the local prevalence and difference in monitoring activity levels of blocklisted entities in these different geographical settings.*

Users on US-WC experience aggressive monitoring activity. Analyzing information depicted in Fig. 2 and Fig. 4a to c, we observe that users located on US-WC run into a smaller number of distinct blocklisted IPs but at the same time register a larger number of hits to these ranges, a clear indication of heightened monitoring activity vis-a-vis other geographical locations.

Nodes located in Europe consistently registered a lower number of blocklisted IP hits when compared to nodes located in Asia. We attempt to maintain a balance while conducting experiments and deploy our code on nearly the same number of nodes in different geographical settings, log data during synchronized time periods. The only difference while gathering measurements in these settings was that we used different lists of queries which were injected into the P2P

network for nodes located in separate continents. For nodes located in Europe we constructed query lists based on European 50 hits [29] and for nodes in Asia we constructed query lists based on Asian hits [36]. The magnitude of difference observed between nodes in Europe and Asia was found to be more or less consistent across the different types of blocklisted IPs. They were however significantly different from measurements gathered across the mainland US. We believe that this difference could again be due to dissimilarity in user behavior and monitoring activity across geographical boundaries.

6 Role Dependent Monitoring

This section delves into whether according to popular perception in P2P communities [10], [11], the probability of being monitored by blocklisted entities varies with the "role" played by a P2P node. The question we answer is: *are UPs monitored with higher probability by entities on blocklists versus regular leaf nodes*. This could show if content providers consider monitoring UPs to be a more fruitful exercise. Through our measurements we find that **there is no conclusive evidence to support any theory regarding role based monitoring**. We observe connection dynamics of UP and leaf nodes in Fig. 5a. Surprisingly, for leaves located in the US the mean number of distinct IPs contacted is higher than for UPs. This is in contrast to nodes in Europe or Asia, where the mean number of distinct IPs contacted is higher for UPs. This observation suggests that *UPs in the US are more conservative in terms of how many users they talk to in comparison with UPs in Europe or Asia*. An obvious question that comes to mind is: should UPs interacting with a lesser number of distinct IPs translate into a lower probability of a UP being monitored? As we will see next this is not always true.

In Fig. 5b we observe the comparison between the percentage of blocklisted IP hits with regards to total IPs contacted for UPs and leaf nodes. This metric depicts if there is any correlation between UPs being monitored preferentially over leaf nodes irrespective of geographical location. We find that UPs in US-WC encounter higher percentages of blocklisted IPs versus leaf nodes. This trend is

Fig. 5. UP Vs Leaf: The black bar signifies UP while the yellow bar signifies leaf users (a) Comparison of average number of distinct IPs contacted by UPs and leaves. (b) Comparison of percentage of blocklisted IPs as encountered by UPs and leaf users.

consistent with Europe based nodes. However for US-EC and Asia based nodes we observe that UPs encounter lesser percentages of blocklist IPs compared to leaf nodes. In fact, we find less than 5% variation in the average number of blocklisted IP hits registered by UPs versus leaf nodes. Thereby we don't find any conclusive evidence for claims of UPs being preferentially monitored by blocklisted entities versus leaf nodes. Also, to answer the question posed previously. Consider the case of US-WC, where UPs talk to less distinct IPs but still are monitored by a larger number of blocklisted IPs. This is clear indication that *monitoring activity varies with geographical location* and that talking to lesser number of IPs doesn't translate into a lesser probability of being monitored. We must mention that our measurements suggest a definite disparity in monitoring activity between US-WC and US-EC and this could possibly be associated to differences in user activity levels at these locations. An imbalance in observations for Europe and Asia can possibly be explained by the "interest" of content providers in trying to monitor P2P networks in those regions. The scanty number of lawsuits in Asia in comparison to significant numbers in the US and Europe provide credence to this explanation [22], [20].

7 Conclusion

To the best of our knowledge, this work is the first to quantify the probability that a user will be monitored i.e. interact with a suspicious IP address. Using Planetlab, we conduct large-scale active measurements, spanning a period of 90 days, from January to March 2006, spread over 3 continents, yielding nearly 100 GB of TCP packet header data. **A naive user is practically guaranteed to be monitored:** we observe that 100% of our peers run into blocklisted users. In fact, 12% to 17% of all distinct IPs contacted by a peer are blocklisted ranges. Interestingly, a little caution can have a significant effect: the top five most prevalent blocklisted IPs contribute to nearly 94% of all blocklisted entities we ran into. This information can help users to reduce their chances of being monitored to just about 1%. At the same time, we examine various different dimensions of the users such as the geographical location and the role of the node in the network. We find that the geographical location, unlike the role, seems to affect the probability of encountering blocklisted users. Finally we answer, who owns blocklisted IP addresses. Interestingly, we find that just 0.5% of all blocklisted IP hits belong explicitly to media companies. The majority of blocklisted users seem to belong to commercial and government organizations and a sizeable portion of the most popular belong to BOGON ranges.

Our work is the first step in monitoring the new phase of "war" between the content providers and the P2P community. It will be very interesting to continue to monitor the evolution of this conflict. A logical next step is to analyze the accuracy and completeness of the blocklists, and the speed with which a new blocklisted entity is flagged.

References

1. <http://peerguardian.sourceforge.net>
2. E. K. Lua, J. Crowcroft, M. Pias, R. Sharma and S. Lim *A Survey and Comparison of Peer-to-Peer Overlay Network Schemes*, IEEE Comm. Survey, March 2004.
3. <http://news.dmusic.com/article/7509>
4. <http://www.betanews.com/article/MPAASuesUsenetTorrentSearchSites>
5. <http://importance.corante.com/archives/005003.html>
6. <http://www.mp3newswire.net/stories/napster.html>
7. <http://news.com.com/2100-1027-995429.html>
8. <http://sourceforge.net/projects/peerprotect>
9. <http://bluetack.co.uk/blc.php>
10. <http://www.boycott-riaa.com/article/9316>
11. <http://slashdot.org/articles/02/05/25/0324248.shtml>
12. <http://www.planet-lab.org>
13. T. Karagiannis, A. Broido, M. Faloutsos, and kc claffy, *Transport layer identification of P2P traffic*, In ACM Sigcomm IMC'04, 2004.
14. E. Markatos, *Tracing a large-scale peer to peer system: an hour in the life of gnutella*, In 2nd IEEE/ACM Intl. Symp. on Cluster Computing & the Grid, 2002.
15. S. Sen and J. Wang, *Analyzing Peer-to-Peer Traffic Across Large Networks*, In ACM SIGCOMM IMW, 2002.
16. Thomas Karagiannis, Pablo Rodriguez and Dina Papagiannaki, *Should Internet Service Providers Fear Peer-Assisted Content Distribution?*, In IMC'05, Berkeley.
17. Kurt Tutschku, *A measurement-based traffic profile of the edonkey filesharing service*, In PAM'04, Antibes Juan-les-Pins, France, 2004.
18. <http://www.techspot.com/news/16394-record-labels-launch-action-kazaa.html>
19. http://www.mpaa.org/CurrentReleases/2004_12_14_WwdeP2PActions.pdf
20. Valerie Alter, Building Rome in a Day: What Should We Expect from the RIAA?, 56 HASTINGS COMM. & ENT. L.J. 155.
21. Jane Black, The Keys to Ending Music Piracy, BUS. WK., Jan. 27, 2003, <http://www.businessweek.com/bwdaily/dnflash/jan2003/>
22. RIAA Gives Advance Warning to Song-Swappers Before Lawsuits are Filed, <http://www.antimusic.com/news/03/oct/item77.shtml>, 2003.
23. Thomas Karagiannis, Andre Broido, Nevil Brownlee, KC Claffy, Michalis Faloutsos, *Is P2P dying or just hiding*, IEEE Globecom 2004.
24. Chu, J., Labonte, K., and Levine, B. N., *Availability and locality measurements of peer-to-peer file systems*. In Proc. of ITCOM '02.
25. F. Clvenot-Perronnin and P. Nain, *Stochastic Fluid Model for P2P Caching Evaluation*, In Proc. of IEEE WCW 2005.
26. http://azureus.sourceforge.net/plugin_details.php_plugin_safepeer
27. <http://mutella.sourceforge.net/>
28. http://www.billboard.com/bbcom/charts/chart_display.jsp?fThe_Billboard_Hot_100
29. <http://www.mp3hits.com/charts/euro>
30. <http://www.trustyfiles.com>
31. <http://isc.sans.org/diary.php?date=2005-04-11>
32. http://www.winmxworld.com/tutorials/block_the_RIAA.html
33. <http://xeex.com>
34. <http://www.completewhois.com/bogons/index.htm>
35. <http://phoenixlabs.org>
36. <http://www.mtvasia/Onair>