IISRC International Industrial Scientific Research Corporation

International Conference on Industrial Scientific Research Engineering

www.worldconferences.org

Robust Swarm Management for Multiple Torrent Environments

GOLLA RAVI KUMAR[1], J. BALA MURALI KRISHNA[2]

[1]PG Scholar, Dept of CS, AITS, Tirupati, AP, India, Email: gravikumar112@gmail.com.

[2]Asst Prof, Dept of CSE, AITS, Tirupati, AP, India, Email: ajbalan.me@gmail.com.

Abstract: This Thesis includes the design, implementation and evaluation of a multi-access peer-to-peer content distribution prototype based on BitTorrent. The prototype is designed to enhance the current BitTorrent in terms of download duration and traffic locality in an overlapping multi access environment. The thesis begins with related state of the art including future ways of networking as well as current state of peer-to-peer networks, object resolution and mobility management. Experiments with conventional BitTorrent show the importance of early stage peer discoveries by the BitTorrent tracker. We therefore enhanced the peer discovery process of BitTorrent and divided the functions of the tracker hierarchically into two layers. The higher layer global tracker keeps track of local trackers, while the local trackers maintain a list of peers in their area of the network. The peer in our prototype retrieves the content locally from its own network whenever possible with extended BitTorrent signaling. While retrieving the content in an overlapping multi access environment, the mobile peer in the prototype recognizes also a new connection to other networks. It is able to use this connection simultaneously with the previous one. The prototype implementation and evaluation show the feasibility of multi access in the content distribution prototype and present the advantages of employing the multi access prototype instead of conventional BitTorrent in overlapping multi access environments in static and pedestrian mobility scenarios. The inter-network traffic was also reduced to almost zero with the prototype in both scenarios.

Keywords: Future Internet, Information-Centric Networking, Traffic Locality, Peer-to-Peer.

I. INTRODUCTION

The BitTorrent protocol has transformed the distribution of large media files due to it's decentralized, scalable, efficient and robust peer-to-peer (P2P) architecture [1]. Peers share bandwidth to help distribute files of common interest. However, the protocol excludes mechanisms for searching, rating, and associating descriptive metadata to, content. Consequently, to locate high quality content available for download, users often rely on web-based systems that provide both links to content (.torrent files) and associated metadata such as a text description, a thumbnail graphic, a URL to associated information on the web and other useful information. This allows users to search and browse available content before making a selection of what to download. Web-based systems rely on user contributions of both. Torrent files that point to available content and metadata that describes content. They are administered centrally so malicious content or incorrect metadata can be manually removed and the users who posted them excluded. Creating new identities in such systems involves some level of user cost since identities must be created on the web system prior to posting information. Hence such systems cannot entirely stop anti-social behavior but they make it costly because creating new user identities involves some time and effort. A number of BitTorrent clients have, recently, integrated their own search and metadata systems such that users can locate and browse available content conveniently from within the client before downloading (e.g. Vuze1 and Miro2). However, again, these rely on centrally administered servers to store and serve metadata requiring the creation of user accounts.

In this paper we propose a design for fully distributed metadata dissemination and rating system which provides similar functionality to centralized systems. Low quality metadata such as spam or incorrect information is combated through a distributed ranking system based on the sampling of user votes. Our approach is a major step towards a completely decentralized, and self-maintaining, BitTorrent media sharing community. As with centralized systems we do not eliminate the possibility of anti-social behavior but we make it costly and difficult without the need for central servers or administration. The target platform for our design is the Tribler4 media client [13] but the design is generic enough that it should be applicable in other media sharing contexts where decentralized and robust metadata dissemination and rating are required. Tribler implements a non-spoof able distributed peer identity system using a public key infrastructure. This means that all communication between nodes is signed and bound to a known source identity, thus preventing forged or stolen identities. Also, the Tribler client provides local database services allowing state to be maintained over sessions.

II. TERMINOLOGY AND RELATED WORK

A. Terminologies

The terminologies used in the BitTorrent community have yet to be standardized. For clarity of exposition, we first define a series of terms to be used in this paper.

Torrent: A BT torrent is the set of peers cooperating to download the same content using the BitTorrent protocol.

Metainfo file: A metainfo file (torrent file) contains all the information used to download the content of interest, including the number of pieces, SHA-1 hashes, the tracker information etc.

Tracker: A tracker is the only centralized component in a torrent. It is not involved in the actual distribution of the content, but keeps track of all peers currently participating in the torrent, and also collects statistics.

Multitracker: The multitracker configuration extends the metainfo file to enable multiple trackers for one torrent. Should one tracker fail, the others can continue supporting the torrent.

Local cluster: It is a set of local peers that download the same content in a given AS. The size of local clusters indicates the amount of local resources available in the AS.

B. Related Works

There have been numerous studies on the implementation, analysis, and optimization of the BitTorrent system; see surveys in [9]. P2P locality has recently attracted particular attention following the pioneering work of Karagiannis et al. [3]. Based on real traces and simulated torrents, they proposed the concept of locality in peer-to-peer systems and evaluated its benefit. Blond et al. [10] showed through a controlled environment that high locality values (defined by [3]) yield up to two orders of magnitude savings on cross-AS traffic, without any significant impact to the peers' download completion time. Xie et al. [11] further suggested cooperation between peer-to-peer applications and ISPs by a new locality architecture, namely, P4P, which can reduce both the external traffic and the average downloading time Choffnes and Bustamante [12] proposed Ono, a BitTorrent extension that leverages a Content Distribution Network (CDN) infrastructure, which effectively locates peers that are close to each other. Bindal et al. [13] also examined a novel approach to enhance BitTorrent traffic locality, namely, biased neighbor selection. Using this method, a peer chooses the majority, but not all, of its neighbors from peers within the same ISP. Our work extends these studies through an Internet-wide measurement that reveals the global distribution of BitTorrent peers as well as the associated tradeoffs of locality.

In particular, we demonstrate that the effectiveness of a locality mechanism can be limited within individual torrents (on which the previous studies have focused), and it is necessary to explore the locality across the multiple torrents. Guo et al. [6] revealed that more than 85 percent of all peers participate in multiple torrents and noted the peer migration behavior. This migration behavior indicates that some BT peers have the potential to serve others even when they have already left the swarm. They proposed an inter torrent approach through tracker-level collaborations. The main idea is to build a tracker site overlay for tracker level collaboration; the peers migrating between different torrents can then be detected and recovered as potential seeders for the torrents. Dan and Carlsson [14] further investigated how the separated torrents can be merged together to improve the performance of an entire torrent. The measurement from Piatek et al. [15] however found that about 91 percent of peers in any single swarm do not arise in any other swarms. This observation seems to contradict the study in [6]; yet this is mainly due to the difference of their objective as well as their measurement schemes. On other hand, the measurement study by Neglia et al. [16] investigated the availability of BitTorrent system among different tracker configurations.

The popularity and the performance of the multi tracker configuration [17] were discussed. Their study showed that around 35 percent of the torrents enable multi tracker configurations. Pouwelse et al. [18] further discussed the relationship between BT trackers and torrents, and examined the tracker availability across multiple websites, albeit with individual torrents. It is worth noting that, the studies of content bundling [19] also provide useful insights to understand multiple torrents behavior. A pioneering work from Menasche et al. [19] studied the content unavailability problem in the BitTorrent system. This study for the first time proposed a model to analyze the availability and the performance implications of bundling through an extensive measurement. Follow up studies such as [20] and [21] also studied some other aspects for content bundling in BitTorrent systems. Considering the measurement and incentive of BitTorrent system, a recent study from Dhungel et al. [22] examined BitTorrent darknets from macroscopic, medium-scopic and microscopic perspectives and investigated the properties of private BitTorrent sites. The study from Otto et al. [4] presented a comprehensive view of BitTorrent, using data from a representative set of 500,000 users sampled over a two year period, located in 169 countries and 3,150 networks.

This study showed that the BT traffic exhibits significant locality across geography and networks. Compare to this study, our work is more focusing on the peer distribution in different torrents/locations. We find that very few individual torrents are able to form large enough local clusters of peers, and this is generally due to the skewed distribution of torrents' popularity. Fan et al. [23] investigated the fundamental tradeoff between keeping fairness and providing good performance for BitTorrent system. Piatek et al. [24] showed that a "win-win" outcome is unlikely to obtain for the ISPs during the locality; the reason is that reducing interdomain traffic reduces costs for some ISPs, while it also reduces revenue for others. Cuevas et al. [25]

also investigated the maximum transit traffic reduction as well as the “win-win” boundaries across the ISPs. Our work was motivated by these studies; yet we explore the multitracker configuration across multiple torrents simultaneously, providing a seamless and lightweight solution to locality in the real BitTorrent system.

III. PEER SAMPLING SERVICE

We assume each peer has access to a peer sampling service (PSS) which periodically returns a random peer from the entire population of online peers. This allows nodes to discover others and potentially exchange messages with them. There are several ways to implement a PSS in a distributed and robust way. One approach uses gossiping or epidemic protocols as shown in Fig.1. Such approaches maintain a random-like overlay network in which nodes regularly exchange their neighbor lists (or view) with others. Such PSS protocols have been shown to be robust, self-repairing, completely decentralized and scalable to tens of millions of nodes [9]. Our target system, the already deployed Tribler system, implements a variant of Newscast [7] called BuddyCast [13].

```
do forever
  wait Δ
  j ← GetRandomNode()
  ml_i ← Extract(local_db)
  Send ml_i to j
  ml_j ← Receive(j)
  local_db ← Merge(local_db, ml_j)

(a) active thread
```

```
do forever
  ml_i ← receive(*)
  ml_j ← Extract(local_db)
  Send ml_j to i
  local_db ← Merge(local_db, ml_i)

(b) passive thread
```

Fig.1. The push/pull gossip based metadata dissemination protocol. GetRandomNode is supplied by the PSS. The local db stores received moderations. The Extract() function returns the moderations list (ml) sent to other nodes. The Merge() function inserts new moderations into the local db. These operations take account of local node votes and moderation regency criteria.

IV. P2P LOCALITY ACROSS MULTIPLE TORRENT: AN OVERVIEW

We now proceed with a framework design for exploring P2P locality across multiple torrents is illustrated in Fig.2. We particularly focus on the tracker-and-client-based solutions [13], which rely only on modifications to end-system implementations. These locality solutions typically replace the random peer selection by an AS hop count-based metric. Upon a request, the modified tracker sorts all other peers in the torrent in ascending order of their AS hop count to the requesting peer, and then sends the prefix of this sorted list (e.g., first 50 peers) to the requesting peer. The requesting peer would then choose the majority, but not all, of its neighbors from peers within the same ISP. Typically, 35 peers within the same ISP (AS hop count 0) can be returned together with 15 other random peers [13]. For the individual torrent scenario, many neighbor selection approaches have been proposed [13], which could also be applied in the multiple torrent scenarios. The new challenge, however, is the detection of peer migrations

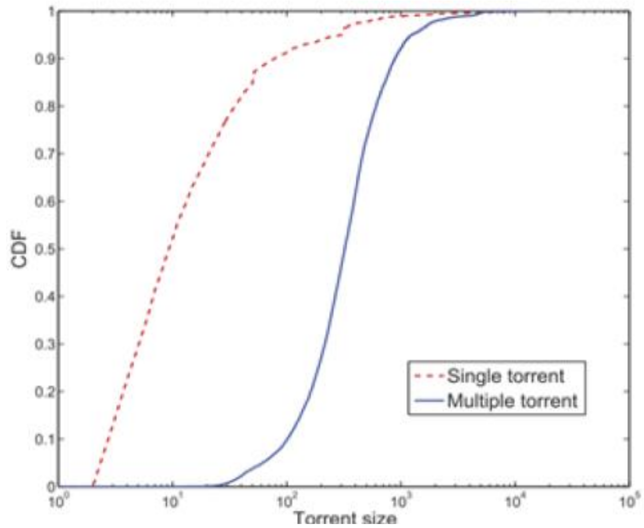

Fig.2. Single torrent versus multiple torrents.

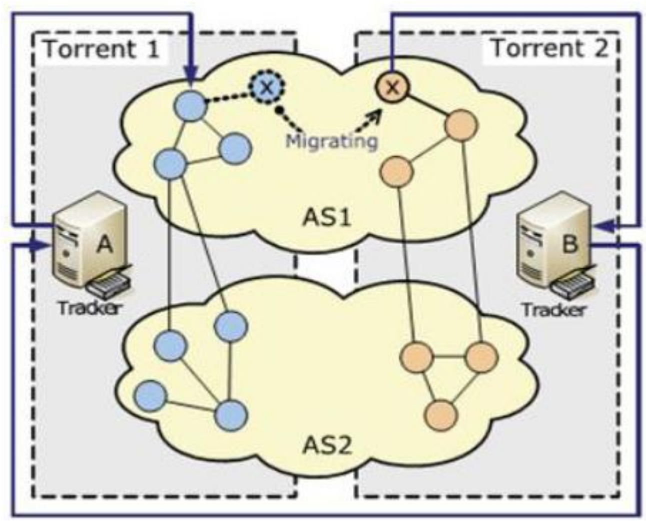

Fig.3. Multiple torrent-based P2P locality.

among torrents. That is, if a peer has finished downloading in a torrent (say Torrent 1) and left, but remains in other torrents, how can we detect it, so as to recover the previously downloaded content to facilitate the locality for the remaining peers in Torrent 1? This is illustrated in Fig.3, where peerxleaves Torrent 1, but remains in Torrent 2. If this migration can be detected, peer xcan still serve as a potential seeder for Torrent 1, which will greatly promote the locality for the peers in AS1. We can see that the solution may need a tracker overlay for tracker-to-peer and tracker-to-tracker communications; in particular, adding extra collaboration among the trackers to trace the migration of peer x[6]. Unfortunately, besides the overheads, enforcing communications between the public trackers can be quite difficult. Table lists site information of the Top 10 most popular trackers in our measurement. We can see that many of them belong to Pirate Bay, which has been involved in a series of lawsuits, as plaintiffs or as defendants. Unless the problem can be well solved, we can hardly expect to organize these public tracker sites together for optimization. We, thus, resort to solutions that minimize the communications, especially tracker-to-tracker communications.

V. RELATED WORK

We first consider the migration detection with shared trackers. Assume Torrents 1 and 2 are both managed by tracker A; any peer migrating between these two torrents can simply be detected by tracker A without communication to other trackers. While this seems to be an ideal case, we now show that it indeed exists and is not uncommon. Our observation starts from the fact that the latest BitTorrent metainfo file can include multiple tracker sites stored in the announce-list section [17]. This multi tracker configuration allows peers to connect to more than one tracker at the same time, which brings two tangible benefits: 1) better accommodates tracker failures, and 2) balances load among the trackers. Fig.5 offers an example with the multi tracker configuration, where Torrent 1 is managed by both trackers A and B, and Torrent 2 is managed both by trackers Band C. In this case, if there is a BT peer x migrating from Torrents 1 to 2, tracker B will receive the arrival message of peer x twice with different content identifications (one arrival message for each torrent). Therefore, tracker B can actually be aware of any peer migration between Torrents 1 and 2 without any tracker level collaboration. The question now becomes 1) how popular is the multi tracker configuration in the real world? and 2) how many migrations can be detected by this configuration in practice? To answer the first question, we consider all the 1,192 trackers in our measurement.

TABLE I: TOP 10 MOST POPULAR TRACKERS

Rank	AS#	Torrents	Tracker Sites (URLs)
1	21202	23386	denis.stalker.h3q.com
2	43350	19915	open.tracker.thepiratebay.org
3	43350	16724	trackeri.rarbg.com
4	21202	15308	tpb.tracker.thepiratebay.org
5	21202	12821	vip.tracker.thepiratebay.org
6	21202	12117	vtv.tracker.thepiratebay.org
7	21202	10019	eztv.tracker.prq.to
8	21202	6079	tracker.prq.to
9	4134	3845	bt1.btally.net
10	15497	3197	inferno.demonoid.com

We record the announce list of the torrents in our data set, and show the cumulative distribution of the trackers that have been used in Fig.4. It indicates that more than 90 percent torrents have specified at least two trackers, and a few torrents even have announced lists of multi hundred trackers. This is much higher than an earlier measurement in 2007 [16] (observed multi trackers in 35 percent of the torrents), and thus suggests the multi tracker configuration has been quickly recognized and deployed in the BitTorrent community. To answer the second question, we model the relationships among different torrents as two matrixes, M1 andM2, where n is the number of torrents in the whole Table 1. Each component of M1, M1i,j 1 is of a binary value, indicating whether torrents i and j have at least one common tracker (1-Yes, 0-No); similarly, each component of M2, M2i,j indicates whether torrents i and j share at least one migrating peer. It is easy to verify that a dot product between these two matrixes, M3 = M1 . M2, gives the detectable migrations by the shared tracker approach.

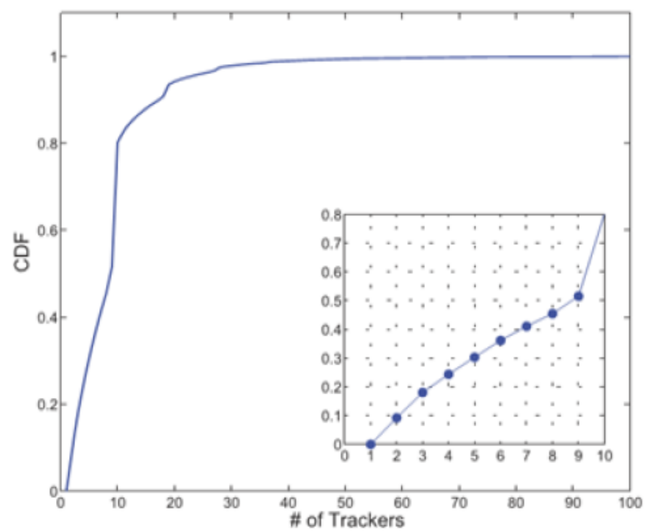

Fig.4. Number of trackers used by torrents.

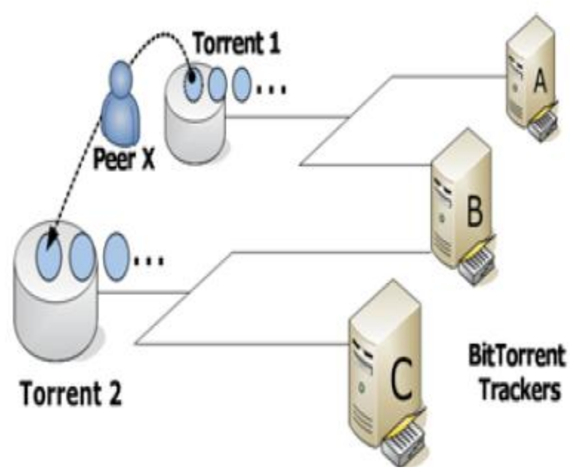

Fig.5. Peer migration in the shared tracker environment.

Specifically, M3i,j=0, indicates that peer migrations between torrents i and j are either undetectable or do not exist at all; otherwise, the migrations between these two torrents will be detected even when M3(i,j)=0. In our measured data, matrix M3 has 2,538 nonzero entities, where M2 has 5,707. Therefore, the peer migrations among about 45 percent torrents can be detected with shared trackers.

Once detected, the shared tracker can then use the biased neighbor selection to improve the P2P locality. It may also forward the migration information to other trackers; however, this collaboration is not compulsory in our framework

VI. CONCLUSION

In this paper, we for the first time investigated the existence and distribution of peer locality across different ASes through a large-scale hybrid Planet Lab-Internet measurement. We found that the BitTorrent peers do exhibit strong geographical locality. However, the effectiveness of a locality mechanism can be quite limited when focusing on individual torrents, given that very few torrents are able to form large enough local clusters. Inspired by the multiple torrent nature of many peers, we proposed a novel framework that traces and extracts the available contents at peers across multiple torrents, thus effectively improving the locality. A series of key design issues were addressed in this framework; in particular, the detection of peer migration across the torrents. Since we can hardly expect to organize the public Internet trackers together for detection, we developed a smart detection mechanism with shared trackers, which incurs no extra communication overhead. It was further enhanced through a torrent clustering approach that explores peer migration patterns.

VII. REFERENCES

[1] CacheLogic, http://www.cachelogic.com/, 2011.

[2]M. Dischinger, A. Mislove, A. Haeberlen, and K.P. Gummadi, "Detect Bittorrent Blocking,"Proc. Eighth ACM SIGCOMM Conf. Internet Measurement (IMC),2008.

[3]T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should Internet Service Providers Fear Peer-Assisted Content Distribution?," Proc. Fifth ACM SIGCOMM Conf. Internet Measurement (IMC),2005.

[4]J.S. Otto, M.A. Sanchez, D.R. Choffnes, F.E. Bustamante, and G. Siganos, "On Blind Mice and the Elephant: Understanding the Network Impact of a Large Distributed System,"Proc. ACM SIGCOMM, 2011.

[5] Planetlab, http://www.planet-lab.org/, 2011.

[6] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "Measurements, Analysis, and Modeling of BitTorrent-Like Systems,"Proc. Fifth ACM SIGCOMM Conf. Internet Measurement (IMC),2005.

[7]The BitTorrent Protocol Specification, http://www. bittorrent. org/beps/bep_0003.html, 2011.

[8] A. Logout, N. Liogkas, E. Kohler, and L. Zhang, "Clustering and Sharing Incentives in BitTorrent Systems," Proc. ACM Conf. Measurement and Modeling of Computer Systems (SIGMATERICS), 2007.

[9] D. Qiu and R. Srikant, "Modeling and Performance Analysis of Bit Torrent-Like Peer-to-Peer Networks," Proc. ACM SIGCOMM, 2004.

[10] S.L. Blond, A. Legout, and W. Dabbous, "Pushing BitTorrent Locality to the Limit," technical report, INRIA, 2008.

[11] H. Xie, R.Y. Yang, A. Krishnamurthy, Y.G. Liu, and A. Silberschatz, "P4P: Provider Portal for Applications," Proc. ACM SIGCOMM, 2008.

[12] D.R. Choffnes and F.E. Bustamante, "Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems,"Proc. ACM SIGCOMM, 2008.

[13] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang, "Improving Traffic Locality in BitTorrent via Biased Neighbor Selection,"Proc. IEEE 26th Int'l Conf. Distributed Computing Systems (ICDCS),2006.

[14] G. Dan and N. Carlsson, "Dynamic Swarm Management for Improved BitTorrent Performance,"Proc. Eighth Int'l Conf. Peer-toPeer Systems (IPTPS), 2009.

[15] M. Piatek, T. Isdal, A. Krishnamurth, and T. Anderson, "One Hop Reputations for Peer to Peer File Sharing Workloads,"Proc. Fifth USENIX Symp. Networked Systems Design and Implementation (NSDI),2008.

[16] G. Neglia, G. Reina, H. Zhang, D. Towsley, A. Venkataramani, and J. Danaher, "Availability in BitTorrent Systems,"Proc. IEEE INFOCOM, 2007.

[17] BitTorrent Multi-Tracker Specification, http://www. bittornado. com/docs/multitracker-spec.txt, 2011.

[18] J.A. Pouwelse, P. Garbacki, D.H.J. Epema, and H.J. Sips, "The Bittorrent P2PFile-Sharing System: Measurements and Analysis," Proc. Fourth Int'l Workshop Peer-to-Peer Systems (IPTPS), 2005.

[19] D. S. Menasche, A. A. A.Rocha, B. Li, D. Towsley, and A. Venkataramani, "Content Availability and Bundling in Swarming Systems,"Proc. Fifth Int'l Conf. Emerging Networking Experiments and Technologies (CoNext),2009.

[20] J. Han, T. Chung, S. Kim, H. Kim, T. Kwon, and Y. Choi, "An Empirical Study on Content Bundling in BitTorrent Swarming System," AsiaFI Summer School, arXiv:1008.2574v1, 2010.

[21] N. Lev-tov, N. Carlsson, Z. Li, C. Williamson, and S. Zhang, "Dynamic File-Selection Policies for Bundling in BitTorrent-Like Systems," Proc. 18th Int'l Workshop Quality of Service (IWQoS), 2010. Wang and Liu: Exploring Peer-To-Peer Locality In Multiple Torrent Environment 1225.

[22] P. Dhungel, D. Wu, Z. Liu, and K. Ross, "BitTorrent Darknets," Proc. IEEE INFOCOM,2010.

[23] B. Fan, J.C. Lui, and D. Chiu, "The Design Tradeoffs of BitTorrent Like File Sharing Protocols," IEEE Trans. Networking, vol. 17, no. 2, pp. 365-376, Apr. 2009.

[24] M. Piatek, H.V. Madhyastha, J.P. John, A. Krishnamurth, and T. Anderson, "Pitfalls for ISP-Friendly P2P Design,"Proc. Eighth ACM Workshop Hot Topics in Networks (HOTNETS), 2009.

[25] R. Cuevas, N. Laoutaris, X. Yang, G. Siganos, and P. Rodriguez, "Deep Diving into BitTorrent Locality,"Proc. IEEE INFOCOM, 2011.

[26] Ctorrent, http://ctorrent.sourceforge.net/, 2011. [27] B. Liu, Y. Cui, Y. Lu, and Y. Xue, "Locality-Awareness in BitTorrent-Like P2P Applications," IEEE Trans. Multimedia, vol. 11, no. 3, pp. 361-371, Apr. 2009.

[28] D. Watts and S. Strogatz, "Collective Dynamics of Small-World Networks,"Nature, vol. 393, no. 6684, p. 409, 1998.

[29] A. Mislove, M. Marcon, K.P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and Analysis of Online

Social Networks,"Proc. Seventh ACM SIGCOMM Conf. Internet Measurement (IMC), 2007.

Author's Profile:

Golla Ravi Kumar, Designation:PG Student(M.Tech-CS), II nd Year, Annamacharya Institute of Technology and Sciences, Tirupati, AP, India. Email Id:-gravikumar112 @gmail.com

J.Bala Murali Krishna, M.E., Designation:-Asst. professor, Dept. of CSE, AITS, Tirupati, AP, Indai. Area of Interest: Computer Networks, Web Technologies and Databases, Data Mining.