A Review on How ISPs do Traffic Engineering of P2P Applications

Rodríguez A.*;Estrada-Jiménez J.*; Tipantuña C.*; Estrada J.**

**Escuela Politécnica Nacional, Facultad de Ingeniería Eléctrica y Electrónica, Quito, Ecuador*
e-mail: {ana.rodriguez; jose.estrada; christian.tipantuna}@epn.edu.ec
*** Corporación Nacional de Telecomunicaciones, Quito, Ecuador*
e-mail: jancoej@gmail.com

Resumen: Los sistemas P2P y los ISPs dependen entre sí para poder ofrecer la mejor experiencia de acceso a contenido al usuario final. Sin embargo, los sistemas P2P no poseen información sobre lo que está pasando en las redes subyacentes, de manera que estos sistemas tienen que cooperar con los ISPs para superar la ineficiencia inherente del enrutamiento P2P. Al parecer esta cooperación tiene que ver con un objetivo principal: localizar tráfico en las redes de los ISPs para aprovechar las características demográficas de distribución de contenido. Este trabajo trata de resumir algunos enfoques, aplicados a las redes P2P más populares, que intentan aprovechar los beneficios de las características demográficas del tráfico P2P, por medio de caching , modificaciones en los protocolos P2P, servicios intermedios gestionados por el ISP, e incluso acuerdos de peering entre ISPs locales.

Palabras clave: P2P; ISP; cooperación; enrutamiento; localidad BitTorrent; Gnutella

Abstract: P2P systems and ISPs depend on each other to give the best accessing content experience for the end user. However, P2P systems have not information about what is happening in the underlay networks, so these systems have to cooperate with ISPs in order to overcome the inherent inefficiency of P2P routing. It seems that this cooperation has to do with a main objective: localizing traffic within the ISPs networks to take advantage of content distribution demographic characteristics. This work summarizes some approaches, applied to the most popular P2P networks, which try to explode the benefits of demographic characteristics of P2P traffic, by means of caching, modifications on the P2P protocols, intermediate services managed by the ISP and even peering agreements between local ISPs.

Keywords: P2P; ISP; cooperation; routing; locality BitTorrent; Gnutella.

1. INTRODUCTION

P2P systems generate more than 50% of the Internet traffic [1], and it sometimes overloads ISPs' infrastructures. P2P traffic is inefficient [2], especially because of its independence on the (underlay) Internet routing. This creates a tension in which the ISP may try to block P2P traffic but P2P systems are able to look for mechanisms to evade this blockage. However, ISPs may have more revenues because of people upgrading their service to use P2P solutions [2] [3], and, if ISPs infrastructure collapses, P2P networks will do so.

It is evident that some kind of cooperation is needed, or any mechanism for the ISPs to be aware of the overlay traffic, or P2P systems to have some information about the underlay conditions, so they can manage traffic in an efficient way.

All the solutions analyzed in this paper take advantage of some kind of cooperation between entities involved in the process of managing P2P traffic. These approaches may be used by ISPs in order to do traffic engineering.

The first approach considers the great potential of locality that a P2P protocol as BitTorrent has, especially because of the characteristics of demand demographics. The second analysis concentrates on optimizing the neighbor selection process of Gnutella, in order to explode locality characteristics of the content shared by P2P systems. The third approach refers to caching P2P traffic as a mechanism to reduce the intensity of P2P traffic on ISP access links. The fourth study is related to the effects of P2P traffic on peering links established between ISPs and the strategies to manage this traffic. The final approach abstracts some of the benefits of solutions intended to efficiently distribute P2P content. Finally some other related solutions and concluding remarks are done.

2. BIT TORRENT LOCALITY

BitTorrent is one of the most used P2P protocols to exchange files in Internet, and it is responsible for more than the 50% of the traffic [1]. Some characteristics of this traffic behaviormay be exploited in obtaining efficiency for the distribution of content.

2.1 Locality

P2P traffic depends, basically, on two elements that ISPs may be able to modify or, at least, to study:

- Demographics of traffic (language, type of content, etc.)

Artículo recibido el 29 de marzo, 2015;.

- Access data rates of users

By taking into account demographics of P2P traffic, locality is the property of this traffic of being significantly concentrated within an ISP [4]. Taking advantage of this inherent characteristic may help to reduce the excessive traffic between Autonomous Systems (AS).

The objective, then, must be to bootstrap P2P users close together in order to maintain P2P traffic within ISPs.

2.2 *BitTorrent random overlay*

A Random selection of neighbors may not always localize traffic well, in particular because of the stratification effect in BitTorrent, which makes peers tend to communicate only with others with similar speeds.

In order to analyze this random strategy, different demographics can be considered where an ISP A may be in different modes regarding to a torrent T [4]:

Sparse mode: when speeds of nodes outside A, participating in T, are very dissimilar from the ones inside A. Because of stratification, internal nodes will communicate only with each other, and the condition for locality will be that a minimum number of nodes must be found as neighbors in the P2P network.

Dense Mode: An ISP is on dense mode when remote nodes have speeds similar to local ones', so stratification is not immediately the answer. A random selection could result in fewer local nodes participating in the P2P network than the ones in sparse mode.

Locality: If a mechanism to keep BitTorrent neighborhood local is applied, when possible, in sparse mode, the probability to localize peers well are higher; however, this locality method used instead of random gets more gains if applied in dense mode, since sparse mode concentrates by itself the traffic inside the ISP.

2.3 *Demographics of BitTorrent*

Due to the way in which BitTorrent works, a crawler can be used to obtain information about indexed torrents, trackers and clients IP addresses, in order to geographically identify the torrents being served. This identification can be done by country or by ISP.

From the study made in [4], most of the torrents are small and the largest torrent has some tens of thousands clients.

Moreover, bounds can be established independently of the peers access speed: an upper bound, considering ISPs on sparse mode (helping to localize traffic); and a lower bound, assuming all torrents are in dense mode. The lower bound, using a Random selection of peers, is almost 0, given that having a random selection of local neighbors in dense mode is a low probability event. In sparse mode, as locality is inherent, the bounds are higher. If a method for locality is considered, bounds get higher.

The effect of speed may be measured based on a metric called inherent localizability, which calculates the density of nodes being part of an ISP in the torrents that it shares with other ISPs that have similar speed. For a Random method of finding neighbors, this metric determines the ability to localize unchokes (a peer answering client requests) whereas on a Locality method, the performance depends on the absolute number of local peers.

The differences in demand demographics of P2P systems are evident, especially taking into account different regions, languages and content development centers, which makes some ISPs the core of download requests, while others become just local centers.

Inside an ISP holding large proportions of content requested by users, inherent localizability is higher, as for example in Spain where searched contents tend to be local and not so global.

Inherent localizability only tends to be increased when local speed is incremented, for ISPs that maintain content locally.

In general, predicting the behavior of inherent localizability is possible by taking into account demographics and speed information.

Global torrents (globally accessed) are present for a wide range of speeds (this means that those are available for download for many different speeds), so increasing the speed inside an ISP will not make its nodes reachable for downloading this content.

However, local torrents are spread at specific ISPs and speeds, so those are as available as global ones. If an ISP increments its speed, reaching other ISPs speeds, some local unchokes may become remote because of the similar speeds. But, if this ISP increases speed further than the others, it will tend to maintain unchokes local.

2.4 *BitTorrent traffic*

Some other parameters like unchoking process, behavior of seeders and optimistic unchokes of leechers, should be considered for a more realistic approach [4].

Seeders normally share its bandwidth based on two criteria: uniform, where nodes of the neighborhood share the same upload rate; and proportional, where the seeder assigns an upload bandwidth proportional to the speed of each node.

Locality is now defined as a set of overlay construction algorithms whose operation is especially dependent on the maximum remote neighbors allowed in a neighborhood N.

Some models of the Locality are:

Locality Only If Faster: Unchokes are switched from remote to local only if local nodes are faster.

Locality: Local nodes are preferred independently of their speed. There is no constrain on the number of remote nodes.

Strict: All switches from remote to local nodes are performed.

2.5 *Impact on ISPs and users*

Implementing Locality may have some drawbacks over the users (i.e. deprivation of users to contact faster remote seeders) but there are bounds where advantages are the same for users and ISPs.

The metrics related to this impact may be: transit traffic reduction compared to random (of interest of the home ISP), and QoS reduction, which is of interest of the user.

When comparing ISPs, those ones hosting global torrents get better savings than those hosting local content, because the first ones have lower inherent localizability. Also, the size of ISPs contributes to this level of saving because bigger ISPs have more possibilities to reach enough local peers.

3. AN ORACLE FOR COOPERATION BETWEEN ISPS AND P2P USERS

This approach concentrates in the neighbor selection process of Gnutella, which is arbitrary done by peers. This is not the best strategy since most of the traffic tends to concentrate in the same AS.

An oracle service implemented by the ISP could help peers to rank the best neighbors, based on its bandwidth, delay and proximity [5]. This rank can be effectively done by ISPs because they have the knowledge of their infrastructure.

The main advantages of using an oracle service are:

- P2P nodes do not have to measure the path performance themselves
- P2P nodes can take advantage of the knowledge of the ISP
- Bottlenecks can be avoided (improved throughput and delay)

After implementing the oracle service, the P2P graph properties are maintained but the ISP has the ability to influence the overlay topology and obtain a biased neighborhood selection. Besides, network performance and delay get improved.

3.1 Evaluation Metrics

The metrics, on which the effect of the biased neighborhood selection depends, are:

- *Degree*: It is the number of outgoing connections. The more connections are the better.
- *Hop count diameter*: The maximum diameter between two peers in a directed overlay graph.
- *AS diameter*: It is the length of the shortest path between the most distanced nodes of a graph.
- *Flow Conductance*: It's a metric related to the failure resistance of the network.

3.2 Overlay/Underlay Graph Properties

The biased neighborhood selection done by the oracle can be compared with an unbiased method where neighbors are chosen randomly. The oracle helps peers to chose neighbors from a list of ranked candidates. The consequences of a biased neighborhood selection are:

- The structural properties does not receive a negative impact since neighbors *remain well connected*, the average path length and the mean degree value do not considerably change
- Traffic is kept within the AS, even when a short candidate list is used.
- Connectivity within AS remains good

3.3 Tests with Gnutella

Gnutella, a file distribution protocol, can be adapted to work with the oracle [5]. The idea consists on sending the Gnutella Hostcache (a list of IP addresses, possible neighbors) to the oracle, before the neighboring is done. The oracle, then, selects its preferred address, based mainly on proximity, and sends it back to the peer. The node will then try to establish a neighbor relationship with the nearest peer.
This kind of neighbor selection does not demean the connectivity of the Gnutella network and the network structure of Gnutella gets unchanged. The result is that peerings are established within the ASs and the percentage of intra-AS P2P traffic gets significantly increased. The scalability of the system gets also increased, since there is much less negotiation messages traversing the network. The general result is that, by consulting the oracle also during the file-exchange stage, and not only for neighborhood selection, leads to a significant increase in locality of P2P traffic.

4. A CACHING SOLUTION USING HPTP

Caching P2P traffic is a more flexible solution, but its implementation is complicated. Some obstacles may be: the lack of standardization in P2P protocols and the economical and administrative cost that the maintenance of these infrastructures represents [6].

4.1 Caching

A caching system intercepts a request from the client and determines if the answer to the request is already in cache (a hit) to directly be delivered. If it is not in cache (a miss), the request is forwarded to the destination. Commonly, the caching servers are located near the gateway of the customers. Besides, significant locality can be found in popular P2P networks like Kazaa and Gnutella, which means that cacheable content do exists.

HPTP has the advantage of using existing caching proxies, so there are no extra adoption costs for ISPs.

4.2 HPTP

HTTP-based-Peer-to-Peer is a framework that has some elements which help to describe the method it uses to cache P2P traffic.

- *HTTPifying tool*. This tool segments P2P streams in small chunks and transports them using HTTP, ensuring that these chunks are cacheable by existing web proxies.
- *Hping.* This tool does cache detection and usability tests. But, locating caching servers is not the only problem; it is also necessary to determine if they can cache HPTP traffic.
  HPing starts by sending requests from a client to the server. Once the server receives this request, a local counter in the server is increased and a response is, then, sent to the client.
  When the client sends an HPing message to the server, this server answers the client, identifying the source address the server sees in the request message. If the address is the same as the client address, there is not a caching proxy but, if the address is not the same, there is a cache.
  The cache usability test, on the other hand, determines if HPTP traffic can be cached by a

specific web proxy system. To do this, HPing sends chained request messages. The destination answers these requests, including the value of a counter which is incremented for each response.

If the counter value is not increased when received at each response, there exists a cache between the nodes.

4.3 Cache Aware Tree Construction (CATC)

The CATC Protocol achieves an explicit control on the selection of caching proxies. The algorithm followed by this protocol is described below:

- A cluster is formed by all peers with the source as its head
- All peers perform cache detection and usability tests against the head
- Peers locally record information about the head
- Peers report their results to a DHT (Dynamic Hash Table) node and they are clustered according to their detected caches
- The DHT nodes appoint the peer whose IP address is the closest one to the source as the new head and inform all peers in the cluster
- These steps are repeated until no new caches are found
- The tree is constructed starting from the best clusters, and within the clusters starting at the head

Handling peer dynamics with HPTP is easier because of the following reasons:

- If a peer leaves or fails, the system is not really affected because the content is already cached
- When a new peer joins, it first tries to find the best caching cluster

Due to the HPTP robustness, it is not necessary to invest on costly infrastructure or optimization, and the maintenance is very simple.

4.4 Remarks on HPTP

The HPTP protocol allows a significant reduction of the backbone traffic since it achieves locality awareness and works better than a locality-aware protocol in the application level.

The number of detected-usable web caches has a linear impact on HPTP performance and the location of caches is especially important when the number of caches is small.

5. P2P TRAFFIC BETWEEN ISPS

Peering relationships between ISPs may cause a unfair distribution of P2P traffic. A relationship could permit an ISP to provide Internet access service to other ISP where an expense is involved for the latter. Free peering (peer-to-peer communication) is also possible between 2 ISPs on the same level, which helps to reduce the dependence on higher level providers. Given the selfishness of P2P applications when using resources to transport information, an ISP A can be affected by other ISP B whose peers use nodes from A to download P2P packets, without paying for the involved transit traffic [7]. Traffic parity between ISPs can be achieved by controlling the links towards the provider and, especially, the free-peering capacity.

5.1 Models of P2P traffic

Let's suppose having two contending local ISPs (1 and 2) providing Internet service by means of a higher level ISP. The link used to connect each local ISP to the Internet is called the private link. The link locally connecting each ISP together is called the peering link. Commonly, only local traffic is allowed through the peering link. Remote traffic (the one generated through the higher-level ISP) can be classified based on its type of demand: private or public (the one needed by the 2 ISPs). Routing of P2P applications in order to predict public demand may also be modeled.

Public demand: An ISP may download content by means of the private link it maintains with a higher level provider, or through a peering link, assuming the aforementioned content was previously downloaded through the private link of the other ISP.

The routing solution: The routing model predicts the aggregated behavior of the routing in the application layer for different traffic intensities, and ISP peering and provisioning agreements.

In the process of optimizing P2P traffic, the following scenarios are possible:

- If cooperation exists between ISPs, a global optimal result is possible if capacities get adjusted
- Having an unlimited peering link capacity, all public demand could be served through only one private link
- If there is no peering, no traffic traverses the peering link
- ISPs cannot efficiently share the public demand because of a peering link without enough capacity

Methodology for ISP economic analysis: The economic analysis of a peering-link strategy can be done by assuming that each ISP has to pay a cost dependent on the intensity of the traffic required to serve all the subscribers. In order to have profits, each ISP will charge its subscribers with a fixed price for a flat rate. Assuming a net income of 0, this fixed price is calculated, and it is called the ISP's break-even price. The benefits of peering strategy will be detected if the break-even price decreases when peering is applied.

5.2 Peering between 2 ISPs

No peering: If there is no local peering, there is economy of scale. ISPs with the highest market share (more caching) are more efficient in using its private link to serve its subscribers. If there is only application routed traffic (P2P public demand), local peering allows ISPs to be as efficient or more efficient than before applying local peering. These benefits may appear in different proportions, even being unfair for some ISPs.

Unlimited-capacity peering: This leads to a complete sharing of P2P downloaded objects between ISPs. Local networks at ISPs tend to behave as one network. Three possible outcomes can be derived from unlimited peering:

- A dominant ISP maintains a lower break-even price and gets even more benefits from peering than the other ISP
- A dominant ISP maintains a lower break-even price but the other one gets more benefits from peering
- A dominant ISP gets less benefits from peering and has a higher break-even price

This scenario may allow an ISP to take more advantage from the peering link by decreasing its private capacity or its traffic towards the other ISP, through the peering link. Obviously, one solution is reducing the peer link traffic in order to limit the usage of the private link from other ISPs.

Limited Capacity Peering: Large ISPs are interested in limiting the capacity of the peering link because it is easier for providers with higher market share to maintain economic advantages. However, this option prevents ISPs depend on the private capacity of other ISPs.

6. IMPACT OF PEER-ASSISTED CONTENT DISTRIBUTION ON ISPS

As it is deeply analyzed in [8], the emergence of P2P applications has helped to provide scalability to content distribution systems. ISPs, however, have received the impact of the increasing traffic generated by peer-assisted solutions. This impact can be analyzed from the perspective of the content provider, the ISPs and the end users. A cache solution, such as the one discussed in a previous section, can alleviate this effect, but its implementation is complex, since it has to be compatible with P2P protocols. In the other hand, peer-assisted distribution algorithms can be adapted to mimic caching systems so they can offer locality-aware solutions that considerably reduce the load on ISPs

Some cooperation is indeed involved in the way that P2P systems work. However, this cooperation does not have to see with ISPs but with the relation among peers. BitTorrent is a classic example of peering-assisted systems where a tit-for-tat policy is implemented by the choke/unchoke mechanism, encouraging cooperation between peers but also banning free-riders. Additionally, potential savings and file availability are inherent characteristics of the demand demographics of these systems. Potential of locality-aware strategies in BitTorrent may be identified by calculating the frequency of users downloading the same content.

Peer overlapping (simultaneous active users) may also provoke a locality effect if at least two peers are downloading the same file. Studies made on BitTorrent protocol showed that approximately 60% of the users could cooperate in the distribution of content as they are overlapping in time. User overlapping is present approximately for 50% of the files, where locality could be applied.

Potential savings of bandwidth can be calculated by determining the bytes unnecessarily downloaded. The experience says that a minimal portion of bytes are downloaded locally. If caching is applied, at least 40% of the content could be downloaded locally.

6.1 Performance for the user

From the user point of view, some aspects can be mentioned about the impact of P2P systems [8]. At least a tenth of the files shared through P2P applications are downloaded more than once within an ISP in a day. These files are short-lived. One third of the users could cooperate for locality and at least 40% of the downloaded content exists locally in active and non-active users.

6.2 Impact of Peer-Assisted content distribution on ISPs

The impact of P2P traffic can be measured in an environment where a big file is downloaded through the BitTorrent network. A locality mechanism within an ISP helps to group IP addresses in order to first look for content inside the ISP. An evaluation done in terms of traffic volume finds that the downloaded traffic can be saved in a factor of two if a peer-assisted locality aware system is used compared with a client-server solution. Also, the uploaded traffic could be reduced by a factor of 6 if locality is applied, because traffic is maintained inside the ISP. And, as it is expected, benefits of locality-aware P2P systems get higher if the size of the ISP is higher.

6.3 Other solutions with locality algorithms

Some other solutions are being studied in order to optimize P2P traffic and maintain it local.

Caching done by an ISP is an interesting option, however, the implementation needs to be compatible with the variety of P2P applications. Besides, a locality-aware solution generates less overhead. Some proxy trackers may also be deployed by ISPs in the network edges, in order to redirect P2P requests to local active users. Other solutions with no extra infrastructure could be implemented by modifying Bit torrent trackers.

Locality solutions based on DNS are not as efficient as a proxy tracker solution located at ISP edges (clients in the same domain span multiple ASs). Matching users by grouping them within a prefix behaves slightly worse than DNS based solutions. A hierarchical-prefix matching algorithm shows the best performance because it dynamically accommodates ISPs of different sizes.

7. OTHER SOLUTIONS

New systems like ONO[9] and P4P [10] have emerged trying to implement locality, by applying changes not only on the client but also cooperatively with the ISP. Previous studies do not take into account the demand demographics or the quantity of peers. P4P is a new paradigm which proposes proactive communications among network providers to basically improve their network resources performance. ONO is a service that supports P2P clients to localize neighboring peers so that traffic among ISPs can be reduced and concentrated locally. As can be seen, there exist solutions which are already trying to modify the current P2P architecture or posing new foundations in order to reduce the impact of applications on the traffic transported by ISPs.

There is also work related to theoretic-games analysis to study economic principles and technical architectures, deriving strategies for ISP to compete inside a local market. Some other studies conclude that uncoordinated overlay and not-overlay traffic would degrade both. Some studies do not take into account the fact that cooperation exists only if two users are active and sharing the same file, as it happens in BitTorrent.

8. CONCLUSIONS

BitTorrent is a P2P protocol which is massively used to exchange big files in Internet. It improves, for example, the downloading experience by distributing the access to the files along a swarm of servers.Although it implements some mechanisms to locallyconcentrate the P2P interactions, more effort is needed to reduce the traffic that P2P applications generate over inter-ISP links due to P2P applications.

Since P2P systems are agnostic of the underlay infrastructure of ISPs and vice versa, cooperation is needed. One of the biggest problems to implement caching is compatibility of different solutions with all the P2P applications. Peering relationships between ISPs should be strictly limited in order to be fair. There are, however, more benefits of locality if the ISP has more users. The benefits of locality may be maintained longer if the selfish behavior of P2P users changes. Locality of P2P traffic influenced by the ISP restores ISP capability of doing traffic engineering.

REFERENCES

[1] Karagiannis, T., Broido, A., Brownlee, N., Claffy, K., & Faloutsos, M. (2004, November). Is p2p dying or just hiding?[p2p traffic measurement]. In Global Telecommunications Conference, 2004. GLOBECOM'04. IEEE (Vol. 3, pp. 1532-1538). IEEE.

[2] Halme, A. (2005). Peer-to-peer traffic: impact on isps and evaluation of traffic management tools. In HUT T-110.551 seminar on internetworking.

[3] Seibert, J., Torres, R., Mellia, M., Munafo, M. M., Nita-Rotaru, C., & Rao, S. (2012). The internet-wide impact of p2p traffic localization on isp profitability. Networking, IEEE/ACM Transactions on, 20(6), 1910-1923.

[4] Rumin, R. C., Laoutaris, N., Yang, X., Siganos, G., & Rodriguez, P. (2011, April). Deep diving into bittorrent locality. In INFOCOM, 2011 Proceedings IEEE (pp. 963-971). IEEE.

[5] Aggarwal, V., Feldmann, A., & Scheideler, C. (2007). Can ISPs and P2P users cooperate for improved performance?. ACM SIGCOMM Computer Communication Review, 37(3), 29-40.

[6] Shen, G., Wang, Y., Xiong, Y., Zhao, B. Y., & Zhang, Z. L. (2007, February). HPTP: Relieving the Tension between ISPs and P2P. In IPTPS.

[7] Hui Wang, J., Chiu, D. M., & Lui, J. C. (2006). Modeling the peering and routing tussle between isps and p2p applications. In Quality of Service, 2006. IWQoS 2006. 14th IEEE International Workshop on (pp. 51-59). IEEE.

[8] Karagiannis, T., Rodriguez, P., & Papagiannaki, K. (2005, October). Should internet service providers fear peer-assisted content distribution?. In Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement (pp. 6-6). USENIX Association.

[9] Choffnes, D. R., & Bustamante, F. E. (2008). Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems. ACM SIGCOMM Computer Communication Review, 38(4), 363-374.

[10] Xie, H., Yang, Y. R., Krishnamurthy, A., Liu, Y. G., & Silberschatz, A. (2008, August). P4P: Provider portal for applications. In ACM SIGCOMM Computer Communication Review (Vol. 38, No. 4, pp. 351-362). ACM.